

This is a section of [doi:10.7551/mitpress/14908.001.0001](https://doi.org/10.7551/mitpress/14908.001.0001)

Principles of Knowledge Auditing

Foundations for Knowledge Management Implementation

By: Patrick Lambe

Citation:

Principles of Knowledge Auditing: Foundations for Knowledge Management Implementation

By: Patrick Lambe

DOI: 10.7551/mitpress/14908.001.0001


ISBN (electronic): 9780262373166

Publisher: The MIT Press

Published: 2023

OA Funding Provided By:

OA Funding from MIT Press Direct to Open



The MIT Press

16 Unhelpful Dualisms: The Personal-Collective Dualism

There is a broad expanse of uncharted territory between the real knowledge work that occurs in an organization and the formal organizational structure and espoused practices.

—Linger et al. (2005, p. 76)

The distinction between personal/individual knowledge and collective or organizational knowledge is a very common one. It implies that an understanding of the knowledge that individuals possess will lead to an understanding of the knowledge that an organization possesses, expressed as organizational capabilities, although, as we have already seen, the mechanisms by which personal knowledge contributes to or is transformed into organizational knowledge are held to be deeply mysterious. Superficially, the implication makes sense. The organization is supposed to harness and leverage the knowledge deployed by individuals.

There has been some debate about whether organizational or collective knowledge is the sum of all individual knowledge, or whether collective knowledge has additional properties that cannot be reduced to the sum of individual contributions (Nelson & Winter, 1982, pp. 104–105; Simon, 1991; Spender, 1996b; Leonard & Sensiper, 1998, p. 122; Chua, 2002; Hardin, 2009, chap. 6; Crane, 2016, pp. 64–68).

However, my view is that organization capabilities do not behave like aggregates of individual knowledges: “Individual interactions are not simply additive, but can take on complex forms and lead to surprising aggregate and emergent outcomes that are hard to predict based on knowledge of the constituent parts” (Barney & Felin, 2013, p. 141).

And as we saw, the consideration of organizational knowledge as an “asset” brings problematic considerations of whether all forms of organizational knowledge can be truly alienated from individuals to the extent that they can be monetized independently of them (Mentzas et al., 2003, p. 137).

The Missing Middle: The Problem with the Personal-Collective Dualism

The dualism between individual knowledge and organizational knowledge is distracting. Personal knowledge and organizational knowledge are only two ends of a complex knowledge integration system (Simon, 1991; Grant, 2002), and neither end makes sense without considering the “*missing middle*”—the spaces where knowledge is constructed, deployed, and consumed in the enactment of work practices within work groups and teams and that are themselves organized around business functions. As Henry Linger et al. (2005) stated:

The concentration on formal organizational programs aimed at the individual workers ignored the real nature of work practices that reside in a space between the organization and the individual. . . . Until the full extent of . . . work practices is articulated, they will remain hidden from the organizational landscape, unappreciated and undervalued. . . . There is a broad expanse of uncharted territory between the real knowledge work that occurs in an organization and the formal organizational structure and espoused practices. (pp. 72–76)

Now, a review of strategic organizational knowledge is all well and good, but it expresses high-level capabilities that resist detailed breakdown, and it tends to focus on future work directions. Such a review it is likely to be much better at reflecting needs, gaps, and goals than it is at inventorying current resources and knowledge in action.

A review of the knowledge held by individuals, particularly key individuals, is all well and good, but the knowledge of everyone in your organization certainly represents a surplus over the knowledge that is actually deployed in key and critical activities. Nor does all the knowledge available in people exactly correspond to what is needed for those activities. There are typically both surpluses and gaps.

In the “*missing middle*”, we have the invisible but crucial knowledge that emerges from people working together to address common tasks and challenges from day to day. Inventorying both strategic and personal knowledge is interesting and has its benefits for specific purposes, but for the purposes of an organizational knowledge audit, it is the knowledge consumed and produced in key operational activities that tells us most of what we need to know.

A more important point is that significant differences exist between the knowledge behaviors of

- individuals,
- teams and work groups, and
- organizations.

These differences are governed by the relative complexity of knowledge integration or knowledge articulation activities at these different levels of scale.

Knowledge integration refers to the ability to bring different forms and parts of knowledge to bear on a problem or an activity. This involves a delicate trade-off between (Grant, 1996; 2002, p. 138)

- (a) supporting specialized knowledge creation and use (i.e., splitting knowledge up among different people) and
- (b) efficiency in knowledge deployment in support of complex tasks (i.e., bringing together different forms of specialized knowledge held by different people).

Elsewhere, this is described as knowledge resource orchestration or resource mobilization, “according to which mobilized resources are integrated into a robust system to support better alignment, coordination and direction for specific use” (Asiaei et al., 2021, p. 1948).

Knowledge articulation is the work we do in groups and organizations that allows us to articulate our knowledge and coordinate our actions effectively without having to discuss and agree on every step in detail (Lambe, 2007, p. 54). “Articulation work names the continuous efforts required in order to bring together the discontinuous elements—of organizations, of professional practices, of technologies—into working configurations” (Suchman, 1996, p. 407).

At the level of the individual, most knowledge integration is done internally within the brain and body, with some external calls on the knowledge of team members or key network contacts. In consequence, a great deal of personal knowledge integration is not easily observed or described. But it can be extremely adaptive and responsive to changes in context.

In teams and in organizations, knowledge integration and knowledge articulation represent activities that need to take place for this scattered knowledge to become actioned effectively. These integration and articulation activities are different for teams and organizations because of the different levels of complexity involved (Kogut & Zander, 1997, p. 312–314).

This is why the dualism between personal and collective knowledge (and as we will see later, the dualism between tacit and explicit knowledge) is misleading. For knowledge managers to be able to design helpful interventions, it is not enough to talk about the knowledge itself. We also need to know the kinds of activities and processes that facilitate knowledge creation, sharing, and use at each of the three levels of individual, group, and organization, and how they differ across levels of complexity. In particular, we need to understand what it takes to facilitate knowledge integration and articulation for teams, as well as the higher-order needs of organizations (Halbwirth & Olsson, 2007, p. 69).

The Special Characteristics of Team Knowledge

Teams and work groups—if they are effective—often have highly developed metacognitive strategies and routines to support interpersonal knowledge integration and articulation (Klein, 1998, pp. 233–257). *Metacognitive strategies* refers to the various ways we become aware of, manage, and organize our own thinking.

In simple terms, this means that members of teams have to develop competencies in how to build and integrate different personal knowledges across the team, if they are going to ensure that the team is coordinated, directed, and kept aligned. Many of these are informal but habituated routines, and many of them are emergent and constructed on the fly, as they are needed. It is a mark of a high-performing team when these metacognitive strategies and routines become largely automatic. The work group's ability to respond to changes in its situation depends on the degree of sophistication of these metacognitive strategies.

From an auditor's point of view, team knowledge has an advantage over personal knowledge. Because this knowledge articulation work depends almost entirely on interactions and the sharing of knowledge artifacts, a team's knowledge use is significantly more observable than individual knowledge use in isolation. Except in highly individualized work contexts, it is also more directly connected to the way that work is done.

Consider the example below from a workshop I conducted with two of my colleagues and a client a few years ago. The italicized portions highlight the metacognitive work done by myself, by my colleagues, and by our client

- (a) to ensure that we are able to integrate the various insights, observations, and learnings as the workshop progresses; and
- (b) to ensure that we are able to articulate our collective knowledge and capabilities in pursuit of common, agreed-upon goals.

Case Study: Team Knowledge in Action

I and two of my colleagues are facilitating a knowledge-mapping workshop with sixteen departments, *about twice the normal size we would normally take on* in one session. It is a global pharmaceutical company with very smart, highly educated, and very vocal participants. *We have already discussed* how we will tackle this large and challenging group.

My colleague Edgar *notices* that one group of participants is getting drawn into very abstract discussions (they are in research and development) and they are not getting their knowledge maps documented, *so he sits with them* to guide them through the mapping process. I *notice* that he is out of circulation, *so I widen my monitoring activity* to include the groups he is no longer covering.

My colleague Jules *notices* that one of the groups is being dominated by a strong, negative character, who “doesn't see the point” of why he should map the knowledge his group uses and

produces. After trying to engage him, Jules *lets me know* there may be a problem with this group. Looking at the guy, *I can see* he is demotivated and unhappy.

My client *has already let me know* that a few strong personalities in the group need to be convinced of the importance of this activity, and *she plans to enlist* her director's support to "lean on" them during the tea break. *Reflecting on the big picture, I know* that the negativity is not universal because *I can see that* most groups are working away well, and Edgar *has told me* that at least one group leader has said the technique is extremely useful and that she wants to take it back and apply it in her own office. However, *I can foresee that pockets of discontent could easily spread* if they are not addressed.

During the tea break (while the "dissenter" is being leaned on), *I tell my colleagues that I want to change the sequence of planned activities* and do a "pain points" diagnostics exercise immediately after the break, before recommencing work on the knowledge maps. We usually use this exercise to consolidate a sense of the most important issues arising from the knowledge-mapping activity. However, it is also a good way of identifying common issues quickly.

They prepare the materials, and after we acknowledge to the group the feedback we have been receiving, *we run the activity*. The room is a hum of energy and discussion. The activity surfaces many shared pain points around poor knowledge sharing and poor access to knowledge when it has to move between departments, and this shared realization, combined with the influence of the senior director on the naysayers, *refocuses* the group. Several participants say they feel relieved at being able to articulate their pain points and at realizing they are not alone in experiencing them.

We had originally intended for the pain points exercise to help the groups integrate and apply the insights from their knowledge maps once they had been completed. *We were able to reconfigure the activity* to show the value of investing effort in the knowledge mapping when we ran into resistance.

After the workshop we *conduct a lessons-learned activity* with the client and *discuss* whether *we should change the plan* for future workshops to position the pain points exercise before the knowledge-mapping activity. Our client *disagrees*, saying that these participants are strong-minded individuals who will always need convincing, and who would not have sufficient context to apply the insights from the pain points exercise until they had performed at least some knowledge-mapping activity. *We decide* to leave the pain points exercise as a midway point within the knowledge-mapping activity to help participants understand how the maps can contribute. However, in future exercises we now have a different mode of running the activity within our *repertoire*.

For myself and my colleagues, this metacognitive work is highly routinized, largely informal, and almost automatic, borne of long experience in facilitating workshops together.

We have a repertoire of skills, practices, routines, and cognitive actions in common that are specific to the team: checking, monitoring, noticing, anticipating, and knowing when to alert colleagues and discuss issues.

It is in this sense that we can speak of team or work-group knowledge that exceeds the sum of individual stocks of personal knowledge. Any new team member would have

to acquire this knowledge from the team as a whole, quite likely modifying the team's collective knowledge in the process.

With the client, we do not have this shared repertoire, so we have to schedule regular check-ins to ensure we maintain a common picture and can adapt our strategies in an agreed-upon way to adapt to unfolding circumstances.

The highlighted cues in the story above illustrate the types of metacognitive work being done at the team level:

- Managing and maintaining common goals and checking expectations
- Noticing and compensating for mutual gaps in awareness, information, and competency
- Informing others of cues they may have missed and maintaining a common situational picture
- Running and comparing mental simulations on the different ways that events might play out
- Adapting planned activities and routines in the confidence that others will step in automatically to provide support
- Checking for the big picture to avoid being led astray by strong but nonrepresentative signals
- Reflecting on events and forming collective views on future action
- Developing new routines to be added to the portfolio for when they are needed

Now, of course, personal knowledge, skills, experience, and judgment are all being engaged here. What the team's metacognitive work enables, is the scaling up of our individual brains, eyes, ears, and abilities, to act harmoniously and in concert on solving common problems, with minimal coordination costs, mistakes, or misunderstandings. The capability of the team exceeds the capability of any individual within the team, and it exceeds the capability of a bunch of individuals who do not have those metacognitive strategies and routines.

This can sometimes have profound consequences. When Captain Chesley Sullenberger landed his Airbus A320 on the Hudson River in New York that cold wintry day in January 2009, improbably saving all 155 passengers and crew, he had been flying with the same crew on a tiring back-to-back flight schedule for four days. They had known each other only slightly before that.

The transcripts of the cockpit recordings and the calmness and efficiency of the entire crew in the crisis reveal the immense value of a highly tuned team. Take this example from Sullenberger's (2009) account:

As I listened to the [cockpit] recording, I saw clearly that Jeff [Skiles, co-pilot] was doing exactly the right things at exactly the right moments. He knew intuitively that because of our short time remaining before landing and our proximity to the surface, he needed to shift his priorities. Without me asking, he began to call out to me the altitude above the surface and the air-speed. . . . Jeff and I had found ourselves in a crucible, a cacophony of automated warnings, synthetic voices, repetitive chimes, radio calls, traffic alerts and ground proximity warnings. Through it all we had to maintain control of the airplane, analyze the situation, take step-by-step action, and make critical decisions without being distracted or panicking. It sounded as if our world was ending, and yet our crew coordination was beautiful. (pp. 380–381)

Drilling together or working together on shared routines has a “tuning in” effect that exceeds the group’s capability to perform just the steps in the routines together. It creates a synchronizing effect even in very chaotic and unpredictable situations. These teams can become surprisingly adaptive and resilient.

Here is an example from Robert Graves’s (1957) account of officer experiences on the western front in World War I, in which the importance of good drill discipline is being discussed in the officers’ mess:

Suppose a section of men with rifles get isolated from the rest of the company, without an N.C.O. in charge, and meet a machine-gun. Under the stress of danger, this section will have that all-one-body feeling of drill, and obey an imaginary word of command. There may be no communication between its members, but there will be a drill movement, with two men naturally opening fire on the machine-gun while the remainder work round, part on the left flank and part on the right; and the final rush will be simultaneous . . . this war, which is unlikely to open out, and must almost certainly end with the collapse, by “attrition,” of one side or the other, will be won by parade-ground tactics—by the simple drill tactics of small units fighting in limited spaces, and in noise and confusion so great that leadership is quite impossible. (pp. 166–167)

Lest you think that this special “all-one-body” capability is just a simple product of highly proficient individuals working together according to common trained standards, the point is *not* that the drills have predicted the precise scenario the troops (or flight crews) find themselves in. They have not. The shared routines, and the mutual knowledge and expectations built up through the drills, give teams the capability and techniques to adapt and coordinate, almost by instinct, when faced with nightmarish situations in which plans and normal operating structures fail.

In fact, there are many examples of air crashes in which the crew is—quite naturally—overwhelmed by the speed with which events are unfolding, by the confusing signals of the instruments, and by poor team coordination and communications. Where teams lack metacognitive strategies for knowledge integration and articulation, they perform less effectively (Bassellier et al., 2003). Knowledge management (KM) can support team

metacognition by addressing how team metacognition works and by providing the processes and tools to build shared metacognition and mutual familiarity.

What we see of the team mind in the “miracle on the Hudson” is backed up by research on the effects of fatigue on the performance of flight crews, especially when faced with unusual or challenging circumstances. Team metacognition builds over time, and it not only powerfully compensates for the negative effects of fatigue but actually enhances the team’s performance beyond expectations, to a degree that the theory of the “sum of individual knowledge” cannot explain.

Research funded by NASA found, rather counterintuitively, that a fresh crew in which the members are new to each other is often outperformed by a *fatigued* crew that has built team metacognition into the way its members work together (Klein, 1998, p. 219; cf. Helmreich & Foushee, 2010, p. 13). The fact that Sullenberger’s crew had been working together intensively for several days mattered a great deal.

In fact, errors are more likely to occur with freshly composed crews (Hackman, 2002; Helmreich & Foushee, 2010, p. 14). This finding led to recommendations to airlines to plan the scheduling of flight crews specifically to maximize crews’ mutual familiarity and to avoid frequent crew reassignments that disrupt this mutual familiarity (Hackman, 2002, p. 50).

Familiarity in this sense does not just mean casual mutual knowledge. Team metacognition incorporates the shared routines, patterns of communication, and mutual expectations that a team builds up as a shared basis for coordinated action, whether through routine interactions, deliberate rehearsal, or a period of coworking.

In April 1980 a US military mission to rescue fifty-two American hostages being held by the regime in Iran had to be aborted because of poor weather, several helicopter malfunctions, and a lack of situational awareness among the task force on the ground. The inquiry into the mission found that one of the critical failure points was the fact that the rescue task force comprised an entirely new organization assembled purely for the purposes of that mission. The failure to run full rehearsals for the mission had exacerbated this weakness. The gaps in mutual familiarity led to a critical breakdown in the task force’s capability and ability to adapt. The mission fell apart, leading to the deaths of eight American servicemen, an extension of the hostage crisis by nine months, and the catastrophic performance of President Jimmy Carter at the ballot box in his reelection bid later that year (Holloway, 1980, pp. 59–60).

Shared knowledge artifacts also matter. In the less dramatic case of our workshop, we were supported by knowledge artifacts—workshop plans, preworkshop briefings, workshop materials, support materials for the activities, profiles of the participants, postworkshop review, and so on. Sullenberger and his colleagues were backed up by training,

instruments, and protocols. Some of our team's lessons and decisions were documented and built into planning templates if we deemed them worthy of reuse. The Sullenberger crew's experience was reviewed by the National Transportation Safety Board to decide what changes needed to be implemented and which knowledge artifacts and training needed to change.

This is consistent with the way that innovation professor Achim Hecker described the collective knowledge resources of functional groups. Without discussing the meta-cognitive aspects, Hecker (2012) points out that the collective knowledge base of functional teams has three essential elements:

- A shared repertoire
- Knowledge of who knows what
- Shared artifacts

In this sense, the knowledge work of the team has more scaffolding and is more visible than the knowledge work of individuals working alone. It is relatively easy to audit compared to personal knowledge.

However, though the presence of team-based knowledge can be identified (with a suitable typology) in an inventory audit, we need to do more than register its presence. We also need to know *how it plays out in behaviors*. This is why we find it useful to conduct a culture audit, specifically based on narrative techniques, alongside an inventory audit. The following case study illustrates why.

Case Study: The *Grafton* and the *Invercauld*—the Power of Team Knowledge

There is another important factor involved in the ability to activate collective knowledge. This is a recognition of, and commitment to, the welfare of the whole group. Practically speaking, this means the ability to form and sustain shared goals, to work collaboratively toward those goals, and to maintain common ground, as well as develop a measure of care for each other (Hackman, 2002, p. 50).

We might call this mutual solicitude, and an important ingredient of this solicitude is mutual trust (cf. Krackhardt, 1992; Wilson, 1996, pp. 138–139). When trust breaks down, team knowledge is also fractured, pushing people back on their individual knowledge resources, so that team performance also suffers. In those circumstances it is extremely difficult to rebuild both trust and team performance without strong participatory interventions—as our chapter 7 case study about the standoff between the local staff and expatriate leadership team in a global nonprofit illustrates.

In 1864, by sheer coincidence, two ships were wrecked on different parts of the same remote island in the Southern Pacific Ocean. The differing fates of their crews can be taken as a fascinating natural experiment in the workings of collective knowledge.

The *Grafton*, with a crew of five, was shipwrecked on January 2, 1864, at the southern end of Auckland Island, several hundred kilometers south of New Zealand. The *Invercauld* was wrecked in May 1864 on the north end of the island, with a crew of twenty-five.

At the time of their wrecks, both ships had a sick crewmate on board. The *Grafton's* crew got their fellow to shore and nursed him to health. The *Invercauld's* crew abandoned their crewmate when they abandoned the ship. From the start, the crew of the *Grafton* were marked by their common purpose, their willingness to go out on a limb for the benefit of the whole crew, and the willingness of team members to defer to the knowledge and skills of the others, regardless of status.

By contrast, the crew of the *Invercauld*, five times larger and nominally greater in aggregate knowledge and thinking capacity, were marked by a spirit of every man for himself, by internal divisions, by violence, by mutual neglect, and by a rigid adherence to status and hierarchy.

By far the most resourceful member of the *Invercauld's* crew in relation to maintaining adequate access to food was Robert Holding, a gamekeeper's son. This was a useful skill set on a remote island. But as he was a mere seaman, neither his peers nor the ship's officers, being socially conscious, heeded his advice on survival strategies (although they were happy to eat the food he caught).

The crew split up, abandoned their fellows, fought with each other, stole food from one another, and ate from one of the corpses. For lack of discipline and common purpose, all but three perished. The captain, his first mate, and Robert Holding were rescued by a passing ship in May 1865, a year after their ordeal had started. They had as little solidarity on the rescue ship as they did on the island, Holding being relegated to the seamen's quarters.

On the south side of the island, where the food stocks were actually more precarious, the *Grafton's* crew had delayed the abandonment of their ship until the last moment so they could salvage as many supplies and tools as possible. They did not have anybody with the same foraging and trapping experience as the *Invercauld's* Robert Holding, but they worked much more effectively together than the crew marooned on the other side of the island.

They built a cabin and a forge, and they experimented with cooking and eating local plants to maintain a balanced diet and avoid scurvy. They established a routine of study, prayer, work, and entertainment. They even caught and raised pets. Living in such straitened circumstances was not without its stresses and strains, but they were solicitous of each other's strengths and weaknesses. One of the crew destroyed his pack of playing cards because losing at cards put their captain, Thomas Musgrave, in such a foul mood.

Fourteen months after their wreck and living in constant hunger, they resolved to adapt their ship's rescued but fragile dinghy for the perilous ocean voyage to reach New Zealand. Realizing it could not carry all five crew members, three of them set sail in mid-July 1865, eighteen months after their shipwreck, reaching Stewart Island, south of New Zealand, a week later. They traveled on to Invercargill in New Zealand immediately, and the gaunt but indefatigable Captain Musgrave set about raising funds for a rescue expedition for his stranded colleagues. A month later he was back at Auckland Island to rescue them. All five of the *Grafton's* crew survived, and it was not until October of that year that they learned that another shipwrecked crew had been just twenty miles north of them for most of their time there (Druett, 2007).

The Special Characteristics of Organizational Knowledge

The collective knowledge activated by teams is marked by its attention to mutual knowledge and mutual goals, and its high adaptive capability. By contrast, the collective knowledge activated by organizations is marked by routinization and the scale of its impact.

Let us look at an example of this kind of collective knowledge in action. Here is former Delta Force commander Pete Blaber's (2008) account of the pre-positioning of men and supplies on an island in the Persian Gulf, preparatory to the US invasion of Afghanistan in 2001:

As soon as the sun went down and visibility dropped to zero, giant cargo planes, flying without lights, came screeching out of the night sky in perfectly synched intervals of ten minutes. As the wheels touched down, the roaring turbocharged engines changed pitch and braked the behemoth flying machines with physics-defying precision. Each cargo plane would then turn off the main runway without a second to spare before another plane, waiting empty at the opposite end of the runway, would release its brakes, accelerate to full power, and go roaring past in the opposite direction to take off and make room for the next plane to land a few minutes later. With Hollywood special effects-like orderliness, this cycle went on all night long. Once the giant planes were in their parking spot, they'd drop their tail ramps and regurgitate their cargo to the ravenous tongues of tandem two-ton forklifts driven by young men who deftly handled the machines like they were Porsches. The entire choreography was done impervious to the naked eye. . . . One wrong turn or second of inattentiveness by any of the actors involved, and a hundred things could kill any of them in a hundred really ugly ways. A small city grew like a weed in front of our mission-focused eyes. (p. 149)

Blaber's description paints a striking picture of large-scale collective capabilities that cannot be explained solely in terms of aggregations of personal knowledge. No single person knows all there is to know about setting up a forwarding base for a large-scale invasion half the world away, even though a multiplicity of deeply specialized knowledge sets are clearly at play.

Blaber's "choreography" is handled through the systematic division of knowledge and of cognitive labor, the provision at functional level of mutual knowledge (i.e., who knows what and who is equipped to handle which tasks), and the shared routines and procedures that are precoordinated and embedded in drills, training, simulations, procedures, and plans. There is a sense in which

knowledge can be "stored" in team or organizational routines, without even having been explicitly described (e.g. as successful sports teams show us). As long as such people and teams remain accessible, one can say that their knowledge is "memorized" by the organization and available for (re)use. (CEN [European Committee for Standardization], 2004, pt. 1, p. 10)

Nobody knows everything. In fact, in relation to the scale of the enterprise, individuals know remarkably little of the knowledge that is being engaged (Hardin, 2009, p. 122).

Teams know more, collectively, than their members do, by virtue of their overlapping but not entirely duplicated repertoires, their knowledge of each other, and their knowledge of how to work together effectively. But they have only the vaguest sense of the larger picture within which they operate.

Everybody has deep knowledge relating to their own small-scale routines and tasks and a light knowledge of how their knowledge, skills, and responsibilities intersect with their immediate coordination partners (Hecker, 2012; Sloman & Fernbach, 2017).

Now there could be an impression from Blaber's description of the airfield growing like a city that organizational knowledge is simply composed of tightly coordinated functional knowledge. But functional or team-level knowledge is not the same as organizational knowledge.

Take the shared knowledge artifacts upon which teams depend to support their work. Many functional groups will store and organize their knowledge artifacts in arrangements that are convenient and transparent to the group, often around the needs of frequent task completion. But these arrangements appear idiosyncratic and opaque to people outside that work group (Lambe, 2007, pp. 146–148, 222–223).

In fact, optimization for functional working often compromises the ability of the enterprise to get an organization-level view of its knowledge artifacts.

This is the source of the much bemoaned *silo working* we complain about in knowledge management—specialized arrangements that serve a useful function for their work groups but that, in the absence of organizational infrastructure to lift them into visibility, inhibit wider organizational uses of knowledge artifacts (Nichani, 2012).

In fact, the fragmentation and scatter of organizational knowledge is often seen as a major challenge in even being able to plan and conduct knowledge audits, because the scope of the audit is not always clear (Lambe, 2017). In one knowledge audit project we were involved in, our lack of visibility into the organizational structure and functions, caused by the devolved and uncoordinated nature of its operations, meant that approximately twice as much work had to be done to cover the ground compared to original estimates.

In practical terms this means that an organization-level view often has poor visibility into functional-level knowledge resources, even when the focus is on explicit knowledge. Paradoxically, organizations can *do a lot more* than teams or individuals but they *know less*. We will return to this thought later.

Blaber's example from the invasion of Afghanistan shows us that organizational knowledge can be striking in its power and capabilities, far exceeding the knowledge

capabilities of individuals or even of teams (cf. Hardin, 2009, pp. 125–126). But organizational knowledge has its limits as well. It is far less flexible than team and functional knowledge, and that in turn is less flexible than personal knowledge in how it can be disposed.

A large number of organizational capabilities are delivered on automatic, and unreflectingly (Nelson & Winter, 1982, pp. 124–125). This is a product of our limited cognitive capabilities. Alfred North Whitehead (1911) stated this succinctly:

Civilization advances by extending the number of important operations which we can perform without thinking about them. Operations of thought are like cavalry charges in a battle—they are strictly limited in number, they require fresh horses, and must only be made at decisive moments. (p. 61)

So organization-level knowledge has quite different characteristics from team or personal knowledge. They can even be in tension with each other, or in conflict.

In another account, Blaber (2008) describes the key battle of Shahi Khot in March 2002, an almost abortive attempt to wipe out an al-Qaeda enclave hidden in a mountainous valley near the Afghan border with Pakistan. The heavily routinized mission-planning process, and the dogmatic mission planners' dependence on helicopters, contrived to set in motion a plan that both individuals and teams could see was set up for failure, given the weather and topography, but that they were powerless to prevent (p. 259).

The large-scale capabilities delivered by precoordinated organizational knowledge come at a cost: an inability to change a complex plan once it has been set in motion. This has been a persistent theme in military history since at least the First World War. Barbara Tuchman (1962, p. 235) relates how the German Kaiser, upon learning of Britain's intent to enter the war, almost immediately changed his mind on the wisdom of the "go" decision to commence hostilities. However, in the space of minutes since his initial assent, the wheels had already been set in motion, and the decision—and the machinery of invasion—could not be recalled. There can be strong tensions and discontinuities between team knowledge and organizational knowledge.

The mechanisms for the integration and articulation of personal and work-group knowledge into organizational performance are sluggish and sclerotic. They are more standardized and formalized. They cannot rely on the informal, responsive, and flexible "coordination by mutual adjustment" that springs from the mutual familiarity of the team (Thompson, 1967, p. 56; Weick & Sutcliffe, 2001, p. 31). That is why the urgent need to combine agility with scale of impact is the holy grail of modern multinational management (Gerstner, 2002, p. 214).

Robert Grant (2002, pp. 138–139) identified four principal knowledge integration mechanisms that organizations use in support of their goals (and that Ralph Stacey

asserted were mechanisms for storing and accessing organizational knowledge; cf. Stacey, 2000, p. 25):

- *Rules and directives*—impersonal mechanisms of varying degrees of formality and strength including standards, policies, procedures, plans, schedules, and other instruments for coordinating large-scale knowledge-based activity.
- *Sequencing*—process definitions that direct how individuals' contributions are inserted into a sequence of activities and tasks. These can function independently of the individuals' knowledge of the whole process.
- *Routines*—simple sequences that are learned and can be performed in different combinations without significant direction or thought.
- *Group problem-solving and decision-making*—typically meetings, designed to bring individual knowledge(s) to bear on an unusual problem and find consensus on a solution. This mechanism is typically deployed in situations for which the other knowledge integration mechanisms are not sufficient. In fact, a proliferation of meetings is a symptom that the other mechanisms are deficient.

These share some similarities with team-based knowledge integration and articulation mechanisms. Both use routines and group problem-solving and decision-making mechanisms, but they work at significantly different levels of formalization.

At the organizational level, relatively few routines are shared across the organization, and they are generalized, not specialized, compared to the work group's repertoire of routines, shared context, and mutual knowledge.

Also at the organization level, the coordination costs of group problem-solving and decision-making activities are much higher when the different parties lack high levels of mutual familiarity and trust. High coordination costs manifest themselves in lots of seemingly unproductive meetings.

In an effective work group, there are considerable economies of effort in this kind of coordination because we can take a shared context for granted, especially where there is a commitment to shared goals, and a mutual solicitude. Things get settled much more easily. Communications can be very compact: "By compactness I mean that a phrase, word or gesture is packed with meaning—meaning that would generally not be extractable by a layperson, without extra information or explanation. Mutual knowledge of various kinds allows for this compactness" (Johannesen, 2008, pp. 199–200).

At the organizational level, shared context (and certainly mutual solicitude) is much thinner, so there is a lot of context-setting, explaining, and context-maintenance work, as well as commitment-maintenance work, in the absence of which there is regular opportunity for misunderstanding, conflict, misalignment, and error.

Within KM we recognize this in the different ways we manage sharing and learning processes. For example, in the transfer of lessons, Nancy Dixon (2000) distinguishes between *serial transfer* (transfer of lessons within a team from one situation to the next), *near transfer* (transfer of lessons between teams sharing very similar contexts), and *far transfer* (transfer of lessons with marked differences in context), each requiring quite distinct support processes.

Over and above their attempts to maintain consistent capabilities in mobilizing knowledge in support of their goals, organizations also go to the trouble of formalizing and protecting some forms of codified knowledge as organizational assets, whether in the form of intellectual property, trade secrets, or proprietary processes and routines. While individuals and teams also create knowledge artifacts, they are typically creating them not as assets but as aids to thinking or memory for individuals, or as coordination devices for teams (Baber & McMaster, 2016, pp. 39–56, 64).

Table 16.1 summarizes the main differences between individual, team/functional, and organizational knowledge. The major differences center around the goal orientation, how the knowledge is applied, the responsiveness of the actors to change, the observability of how knowledge is being used, the complexity of the knowledge work, and the influence of culture on the use of knowledge.

Organizational knowledge as a set of capabilities and resources is therefore managerially and phenomenologically distinct from the functional knowledge leveraged by teams and workgroups, and this is again quite distinct from personal or individual knowledge.

Does it make sense to speak of “team knowledge” and “organizational knowledge” as if they are “things”? I believe it does. There are those who argue against the “reification” of collective knowledge as if it exists independently of individual knowledge, following Herbert Simon’s (1991) declaration that only individual humans can “know” (p. 125; Grant, 1996, p. 113). I think this is overly restrictive. As Paul Cilliers (2000) put it, human knowledge is constituted socially and cannot subsist outside of social relations:

Knowledge comes to be in a dynamic network of interactions, a network that does not have distinctive borders. On the other hand, this perspective would also deny that knowledge is something purely subjective, mainly because one cannot conceive of the subject as something prior to the “network of knowledge,” but rather as something constituted within that network. (pp. 8–9; cf. Sloman & Fernbach, 2017)

Of course, groups and organizations are not single-brained creatures as humans are; they have many brains that have to be aligned. But on the converse side, humans with their single brains do not possess their knowledge entirely independently of others. Knowledge is a shared resource and is possessed and acted upon both within and

Table 16.1

Differences between individual, work-group, and organizational knowledge

	Knowledge of individuals	Knowledge of teams and work groups	Knowledge of organizations
Goal orientation	Typically focused on task completion.	Focused on outcomes of the integration of multiple tasks. Requires coordination and goal-setting activities and regular check-ins to repair differences in perception about goals and how to achieve them.	Focused on high-level organizational goals and typically communicated in broadcast mode, with poor feedback loops and repair mechanisms for divergent understandings of goals compared to those of the work group.
Application of knowledge	Application of knowledge is largely habituated.	Application of knowledge is routinized in high-performing teams. The team develops both formal and informal metacognitive strategies and routines for building and maintaining common ground, maintaining member situation awareness, goal setting and tracking, and compensating for performance, knowledge, and understanding gaps in members (Klein, 1998, pp. 233–257).	Emergent, loosely supervised, and loosely controlled application of individual and team knowledge.
Responsiveness to change	Emergent, highly responsive to context.	Highly responsive in teams with good metacognitive strategies and routines. Poor metacognitive strategies and routines can produce mistakes and poor responsiveness to changing contexts.	Relatively unresponsive. Improvements in responsiveness carry high coordination costs and may be impeded by culture and politics.
Observability of knowledge work	Hard to observe, high effort required to describe in any detail.	More easily observed than individual knowledge, knowledge activities are made visible through interactions and trading of knowledge artifacts.	Observed and described at high levels of abstraction or, in detail, in limited ways for the purpose of legal protection through contracts, intellectual property, trade secrets, and privacy or confidentiality regulations.
Complexity of knowledge work	Primary focus is on direct knowledge application in support of personal and joint tasks; secondary focus is on formal and informal knowledge integration and articulation.	Primary focus is on informal knowledge integration and articulation in support of joint tasks; secondary focus is on formal knowledge integration and articulation; tertiary focus is on ensuring availability of requisite knowledge.	Primary focus is on formal knowledge integration mechanisms (e.g., policies, procedures, routines); secondary focus is on informal knowledge integration activities. Support for knowledge articulation is typically weak.
Influences	Influenced primarily by team or work-group culture.	Influenced primarily by team or work-group culture.	Influenced by the emergent organization-wide effects of competing subcultures.
Role of knowledge artifacts	Mainly as memory and sensemaking aids, often informal.	Used to maintain situation awareness and collaborative sensemaking, both informal and formal.	Formal, focused on reporting, organizing action, analysis, and record keeping.

between human beings, at differing levels of scale. As long as we realize we are using a “thing”-based label metaphorically and not literally, and recognize that it does not completely express the dynamism of knowledge work, then the benefits of doing so (the entailments of the metaphor that knowledge can be managed, accumulated, built upon, and deployed) outweigh the limitations.

In this sense, the stable, enduring factors that persist as individuals come and go, and that give teams and organizations the capability to extend, apply, and gain knowledge in pursuit of common goals can legitimately be called forms of knowledge. As Gary Klein (1998, p. 257) points out, while the parallels between human minds and team minds should not be taken too literally, they share features in common (e.g., limited working memory, storage of information, limited attention spans, parallel processing of information), and the workings of each can illuminate the workings of the other (cf. Hardin, 2009, p. 121).

A Middle-Out Method for Auditing Knowledge Stocks

What then are the implications for knowledge auditing if we break down the dualism between personal knowledge and organizational knowledge into a tripartite distinction, looking at personal, team, and organizational knowledge?

From an auditing point of view, we can see that functional (work group) knowledge provides an important window into how knowledge is produced and used in the organization. Figure 16.1 shows schematically how functional knowledge mediates




Figure 16.1
The role of functional (work group) knowledge in mediating personal knowledge to the organization.

personal knowledge to the organization and can support what we call a *middle-out method* for knowledge auditing.

First, functional knowledge, enacted within work groups and teams, renders the aspects of personal knowledge that are of interest to the organization's objectives and goals both observable and actionable.

Second, to understand organizational knowledge, it is necessary to understand the particulars of functional knowledge, because this is what the organization has to harness and utilize and coordinate to meet its goals. Starting an analysis with an examination of organizational knowledge is an abstract and theoretical exercise at best, until there is a prior understanding of what tasks are actually performed and how.

This is why the dualism between personal and collective knowledge is too simplistic. The "missing middle" prevents us from understanding how personal knowledge is mediated and rendered actionable to the organization.

On the other hand, to understand the functional knowledge in an organization it is neither necessary to understand all the particulars of personal knowledge nor practical to do so. Personal knowledge typically exceeds that which is required for the performance of organizational tasks.

A good characterization of functional knowledge will indicate which particulars of personal knowledge need to be investigated and described in greater depth. This is why the focus of an inventory audit should be organized primarily around functional knowledge in what Linger et al. (2005, p. 72) call a "middle-out" approach. This is consistent with Nonaka and Takeuchi's (1995, chap. 5) advocacy of *middle-up-down* knowledge creation processes.

This does not mean we should disregard either personal knowledge or organizational knowledge. It simply means that an inventory audit organized around functional knowledge provides a practical avenue to understanding which personal knowledge resources and needs may need further investigation, on the one hand, and which strategic and organizational knowledge resources need attention, on the other hand.

Functional knowledge audits provide the lens through which to identify the most important dependencies on personal knowledge. For example, a functional knowledge audit that revealed a high dependency on certain areas of specialized tacit knowledge would provide the warrant to investigate those areas in greater depth with the relevant individuals.

When we come to reviewing strategic knowledge audits later on, we will see that strategic knowledge needs and capabilities are typically mapped against strategic organizational goals. Strategic knowledge needs do not necessarily match existing functional knowledge

resources, especially where new strategies are being pursued and new capabilities need to be built.

When conducting a strategic knowledge audit, the functional knowledge audit can perform a valuable preparatory function. It enables us to compare what we have at the functional level against what we need at the strategic level. We can look at how functional knowledge (and behind that, in targeted areas, personal knowledge) supports those strategic goals and what must be done to improve that support—that is, where we have the strategic knowledge resources we need, and where we do not.

So although it has not been especially common to think of knowledge as primarily a team-level phenomenon, there are good arguments for doing so. And it turns out that there are some good precedents for inventorying knowledge resources at the function and task level.

Both Karl Wiig and Dave Snowden have used ethnographic techniques to analyze knowledge needs and uses at the task and function level. Wiig followed a traditional ethnographic approach of using skilled interviewers, while Snowden advocated self-ethnography, a facilitated method for self-reporting knowledge use (Wiig, 1995, chap. 8, 10; Spradley, 1979; Crane, 2016, p. xix). Self-ethnography is an example of the participatory approaches we described in chapter 7, where evidence collection and interpretation are driven by the participants themselves.

Wiig focused on the task and the task environment, while Snowden focused on major decision points and decision clusters, assuming that key decisions represent the critical points in tasks where knowledge is most closely engaged and revealed. Both Wiig and Snowden used this analysis to map information and knowledge resources and flows associated with tasks and decisions.

In the information audit domain, Elizabeth Orna (2004) took an implicitly functional approach. While she did not orient her questions toward the particulars of performed tasks, she worked systematically through different functional departments and got them to collectively map their knowledge and information needs, uses, and flows in a workshop format. This approach was also advocated by Dave Snowden and by Max Boisot (1998, p. 231).

Workshop formats using group-mapping techniques are especially good at surfacing shared representations of collective, team-based knowledge, and are superior (for this purpose) to interviews, which are optimized for exploring aspects of personal knowledge. These workshops are classic examples of participatory approaches. Collective self-representations in workshops mitigate against the subjectivity of individual perspectives: “We are thus initially dealing with a collection of individual perspectives

and representations that are subjective in nature and that confront each other. Gains in objectivity are achieved gradually through iterated discussions and further investigations” (Boisot, 1998, p. 231). And in consequence, key tasks and decision points quite naturally emerge as the main organizing elements in the resulting maps (Orna, 2004, pp. 69–71).

The functional and task-oriented approach is one my company Straits Knowledge has used for the past twenty years, adapting features from Boisot’s, Orna’s, and Snowden’s approaches. We facilitate workshops with knowledgeable and experienced representatives of work groups to help them self-report a shared understanding of their critical tasks and activities and the knowledge uses associated with those. The approach is deliberately participatory. In this way we mitigate the risk of bias from partial views, poor self-knowledge, and perspectives dominated by the individual’s knowledge rather than the team’s knowledge.

Case Study: Collectives “See” More than Individuals

Be My Eyes is a Danish tech start-up launched in 2015, with a free iOS application that connects sighted volunteers worldwide with people who are blind or have low vision. An Android version followed in 2017. The idea is ingenious. If you are blind and in a situation in which you need to figure out what a package label says, the nature of the object you are handling, or a direction sign, you open the app and put out a call for help. Public transport and shopping can be particularly challenging situations to navigate for people who are blind or have low vision. When preparing food, you may need help to identify which food cans or packages you need.

Typically within a minute, one of the 1.5 million volunteers signed up with the service will connect with you via video call, you will point your smartphone at the thing you need help with, and the volunteer will act as your “eyes” to describe to you what they see. As of October 2018, over 100,000 people who are blind or low sighted were using the app. There is also a business version that helps companies meet the needs of customers who are blind or low-sighted (Be My Eyes, 2018).

Other start-ups, such as Aira, are more commercially oriented, using wearable smart glasses and streaming video to connect clients who are not sighted with paid, sighted hourly contractors to help them perform tasks. Aira’s plan is to use artificial intelligence to build learning algorithms around repetitive tasks so that it will eventually be able to recognize certain tasks and provide automated assistance (Lee, 2017).

Both examples provide particularly vivid illustrations of a fact that we know implicitly but do not always acknowledge in the way we investigate knowledge use in organizations. Individuals on their own see less and know less than collectives working together. We are all partially sighted to some degree, constrained by our work contexts and by the daily events that direct our attention and interest. Knowledge maps built by individuals are invariably poorer, and commonly contain significant gaps and inconsistencies, compared to maps built by groups working together.

Summary

In this chapter we broke apart the simplistic distinction between personal and organizational knowledge. Here is a summary of the main points:

1. Personal knowledge, team knowledge, and organizational knowledge are phenomenologically and managerially different from each other and are not simply composites or reducible elements of each other.
2. In an inventory audit, it is necessary to observe and understand the metacognitive strategies for integrating knowledge use within and across teams.
3. A knowledge resources inventory needs to document knowledge uses at the functional or team level in a middle-out method for knowledge auditing that can be extended downward to personal knowledge where necessary and mapped upward to strategic needs at the organizational level.
4. Functionally oriented knowledge audits can help to mediate an understanding of how personal knowledge contributes to strategic organizational capabilities.

© 2023 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.
Subject to such license, all rights are reserved.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Lambe, Patrick, 1960– author.

Title: Principles of knowledge auditing : foundations for knowledge management implementation / Patrick Lambe.

Description: Cambridge, Massachusetts : The MIT Press, [2023] | Includes bibliographical references and index.

Identifiers: LCCN 2022015397 (print) | LCCN 2022015398 (ebook) | ISBN 9780262545037 (paperback) | ISBN 9780262373159 (epub) | ISBN 9780262373166 (pdf)

Subjects: LCSH: Knowledge management. | Organizational learning. | Organizational change.

Classification: LCC HD30.2 .L362 2022 (print) | LCC HD30.2 (ebook) | DDC 658.4/038—dc23/eng/20220718

LC record available at <https://lcn.loc.gov/2022015397>

LC ebook record available at <https://lcn.loc.gov/2022015398>