1

THE NATURE OF KNOWLEDGE

1.1 KNOWLEDGE: THE MOST PRECIOUS ASSET AND THE GREATEST CHALLENGE

> Knowledge represents the "…highest value, the most human contribution, the greatest relevance to decisions and actions, and the greatest dependence on a specific situation or context." (Grover and Davenport, 2001: 6)

At the turn of the millennium, Varun Grover and Thomas Davenport reflected the opinions of many involved or interested in the organizational practice of knowledge management (KM) when they cast knowledge in the role of the organization's top prize. It is also one of the most challenging and complex topics on the organizational agenda.

The value of knowledge to the organization is in fact one of the few areas of consensus in a field otherwise defined by its many debates, controversies, and disagreements. With the modern organization operating in an increasingly complex world marked by change and uncertainty, that value can only intensify. Consequently, knowledge has come to be seen as a firm's most precious asset, the key to new product and service development, essential to understanding customers and market trends, and the principle ingredient to innovation, to name just some of its stellar attributes. It is also the asset that most easily walks out of the door.

Yet, as it turns out, the nature of knowledge is one of the greatest challenges facing organizations and the field of KM, according to the economist and business consultant Robert M. Grant and many others. This one word stirs up more of a storm of controversy than any other issue. As Grant points out, the search for a *definition* of knowledge is far

Knowledge and Discourse Matters: Relocating Knowledge Management's Sphere of Interest onto Language, First Edition. Lesley Crane.

from resolved, although some in the academic domain of philosophy might well dispute that.

This chapter has two main objectives. Firstly, it attempts to convince on the salience and importance of engaging with the question over the nature of knowledge. Secondly, it offers an overview of the shape of the debate in KM: what does it look like, and how does the perception of the knowledge phenomenon impact on its practice? The shape of the debate is itself something of a challenge: it does not conform to a dialogic debate in which one opinion can be seen as building on another, but rather is pulled this way and that, splintering off in this or that tangent. The question of the nature of knowledge is so broad that it is very easy to succumb to the superlative or to lose sight of the objective in a tangle of deep thought inspired by philosophical accounts. To be clear, we are interested in a conceptualization of knowledge from the perspective of the organization, and the practice of KM, and not that (or those) conceptualization(s) discussed within philosophy.

So, beginning with an overview of why this question is important, the discussions open up a critical review of the various ways in which knowledge is described with particular interest in what is arguably the most widespread and popular, yet most contested, of these: the tacit–explicit duality. This is contrasted with alternative ways of understanding knowledge, which are shown, in subsequent chapters, to be consistent at least in part with a perspective with considerable empirical support. This leads to a view of knowledge that is developed throughout this book as the central tenet of a way of thinking about knowledge and its management that extends some directions that many in the field have been indicating for the last two decades or more.

1.2 WHY AN UNDERSTANDING OF THE NATURE OF KNOWLEDGE IS CRUCIAL

Given the eons of debate around the nature of knowledge, it is hardly surprising that a consensus eludes the field of KM. For instance, Peter Heisig, a researcher at the University of Cambridge's Engineering Design Centre, trawled through 160 different KM theories, concluding that a "uniform understanding" of knowledge is simply not available. Instead, there is a melting pot of confusion, ambiguity, and contradiction over its nature, constitution, and even its location (see the end of the chapter for a list of suggested further reading).

Of course, the irony in all of this is that while on the one hand knowledge and its management is understood as a vital part of organizational strategy, as the metaphorical key to organizational success and growth, there is on the other the absence of anything resembling a "uniformly" accepted definition. Now allowing for the fact that large disparate groups rarely agree on definitional matters (political parties, for instance, and the question of what constitutes economic success), this state of affairs inevitably casts a shadow of doubt over both theory and practice. Quite simply, if a product cannot be specified, how can it be the subject of theory? If it cannot be grasped, how can it be managed, measured, or otherwise leveraged? How can the manager claim success in managing knowledge, if there is disagreement over what it is? Ask any group of people within an organization to describe what they mean by knowledge, and it is virtually guaranteed to result in as many variations. Michael Polanyi, undoubtedly one of the most influential

voices in the field—albeit indirectly—proposed that definition rises from "formalization" of meaning. From this perspective, a definition of knowledge would represent a generally accepted and unambiguous meaning shared by the majority. And in this case, "majority" refers to those within the KM field.

The bottom line is this: if it is the role of KM to improve the handling of knowledge, then it is important to arrive at an understanding of what knowledge is—that is, an understanding that is relevant to the organization and its activities, goals, and so forth, which can be subscribed to by the majority.

Developing this line of thought, Polanyi proposes that a word's meaning takes shape in its repeated usage. More specifically, Professor and psychologist Kenneth Gergen suggests that the definition of knowledge is no more or less than as expressed in the language conventions favored by certain groups at certain times—that is, "meaning" as truth exists in shared consensus. Or, put another way, it is whatever is fashionable with a majority at any given time.

These viewpoints can be interpreted as diffusing any sense of urgency for definition and even as establishing a rationale for seeing a definition of knowledge as whatever is fashionable. But this raises a question: fashionable for whom, and shared and used by which groups? Is one group's viewpoint more important than another's? Or more right? Also, while the "fashionable perspective" offers the prospect of an escape route from a troublesome debate, it does not help in addressing other questions: Is knowledge objective, subjective, or both? Can knowledge only ever be personal, rooted to the individual's experiences, beliefs, attitudes, and so forth, or can it exist as a communally shared phenomenon? Can knowledge ever be seen as a true representation of the world, or is it entirely socially constructed, influenced, and molded by human action? Is it an object or a commodity—is it a practice, or can it be either? And so on, and so forth. This is the proverbial tin of worms—and this is what KM sets out to manage.

These are complex questions, with profound consequences for the practice of KM. If knowledge is seen as one particular color, with energies and strategies geared around that, then there is a risk of ignoring knowledge of another color. Haridimos Tsoukas of the Athens Laboratory of Business Administration and Efi Vladimirou of Planet Ernst and Young, in their essay on the relationship between organizational and personal knowledge, take an even more pragmatic approach: criticizing those who claim there to be no need for a "concrete definition of knowledge," they reason that where there is theoretical confusion, it should not be a case of simply abandoning theory, but rather it should promote the search for clarity and cohesion.

Arguably, it is this lack of a definition of knowledge that underpins many of the issues and debates, which characterize the field. For example, researchers and authors France Bouthillier and Kathleen Shearer, writing in 2002, point to 18 different definitions of KM practice in their investigations of methodologies adopted by various organizations in published case studies. Almost a decade later, Patrick Lambe, principal consultant at Straits Knowledge in Singapore, also draws attention to this issue, remarking that throughout most of the 1990s, the same names were being applied to quite different concepts by different people in the KM field. The question over what constitutes the practice of KM is taken up in more depth in the following chapter.

The point should be adequately made by now. This is murky territory. So murky, in fact, that many scholars, researchers, and practitioners simply avoid any reference to the nature of what they focus on. For instance, a study of KM practice's impact on firms'

innovation capabilities investigates the human barriers associated with knowledge generation, codifying, sharing, and application, but does not define the subject of the research. Another research study investigates how processes such as knowledge creation and between-firm knowledge sharing impact on innovation in firms in three different countries but also neglects to define what is meant by "knowledge." Even Thomas Davenport, one of the most respected scholars and authors in the field, and who was named one of the top 25 consultants in the world in 2003,[1] is not immune. In his illustration of how technologies can be successfully deployed in firms for the management of knowledge, he seems to step aside from any definition of the subject of his investigations, although offering detailed and thorough discussions around, for instance, types of knowledge companies and attributes of the knowledge manager. He is, however, a prolific and widely respected author and authority and is likely to have tackled this issue elsewhere. Nonetheless, it is perhaps revealing that, in the context of a book aimed at business managers, this particular text avoids the question. Is it simply that no one is interested?

Here is another reason why the nature of knowledge *should* be of interest to the organization's leaders and practitioners: if there is a mountain of viewpoints on what constitutes knowledge, then, as noted earlier, this unavoidably has an impact on the definition of KM itself. And that has consequences for *what* is practiced.

Lambe's reference to the "1990s" suggests phases of development. David Snowden, founder of Cognitive Edge based in England, offers a concise and useful developmental context against which to pin the knowledge definition debates. He proposes three ages or "generations" of development. The first, before 1995, is very much dominated by technology and process reengineering, with its focus on managing the flow of knowledge to decision-makers. This was the age of "human replacement"—if a machine could do the job of the worker, then the worker had to go. The second generation, beginning in the mid-1990s, is ignited by the publication of Ikujiro Nonaka and Hirotaka Takeuchi's groundbreaking and influential book, *The Knowledge-Creating Company*, and its introduction of what would become the most famous—for some, infamous—perspective on knowledge: the tacit–explicit duality. The start of the new millennium ushers in a third generation, mobilized by Snowden himself with his formulation of knowledge as paradoxically both a "thing" and a "flow." As the following discussions illustrate, both notions of knowledge are very much part of the polarized dialogue.

So, if theoretical confusion over the nature of knowledge is the case, the solution is not to avoid the issue—however tempting that might seem—but rather to evolve more and better theory, as argued by Tsoukas and his coworker. So how is "knowledge" defined by those brave enough to take a stand?

1.3 WAYS OF DEFINING KNOWLEDGE AND THE RISE OF A SINGLE PERSPECTIVE

It has been suggested here that the labels applied to knowledge—the *ways* in which it is defined—have a significant impact on its management, both from the perspective of a practice and a subject of research. What we find is that groups of opinions are ranged like troops on a battle field: no matter which opinion looks strong and well equipped, there is

[1] Consulting Magazine.

always an opposing force. Such matters are important because the stakes are high. Writing about knowledge and competitive advantage in 2007, scholars Laurence Prusak and Leigh Weiss point to a report that estimates that an organization employing 1000 knowledge workers could waste as much as US$6 million a year on time spent searching for nonexistent knowledge or duplicating knowledge that already exists. One could, however, question what it was that "employees" are looking for or duplicating—knowledge or information? A more recent study by Dublin Institute of Technology researchers Mohamed Ragab and Amr Arisha equates the loss of one experienced manager from a Fortune 500 company to around US$1 million worth of lost knowledge. An organization's knowledge needs to be managed, but manage what? There are three underlying connected themes at issue: knowledge as an object versus knowledge connected to social action, knowledge as objective or subjective, and personal versus organizational knowledge.

1.3.1 Knowledge: "Thing" or "Action"?

The lines can generally be drawn up into those who—explicitly or otherwise—reify knowledge as an object, and those who view it as social action. On the latter, "social action" refers to the notion of knowledge as action accomplished in interaction with others. This latter understanding requires a little explanation. Taken to its extreme, the idea of knowledge as action accomplished socially might suggest that it is only in social interaction with other people that knowledge comes into existence. But this interpretation causes problems for the idea of individual capability to experience a "light-bulb" moment in solitude, for instance. Alternatively, it could be interpreted as "knowledge as a psychological phenomenon experienced in social interaction with others, or through the individual's interaction with their environment." Where knowledge is approached as constituting social interaction or as its product, this does not necessarily imply that knowledge is always spontaneous to that interaction. In other words, the germ of knowledge new to the recipient may be seeded in social interaction with others, but may not emerge as such until some later point, perhaps triggered by some unconnected event in the environment. It is to this latter understanding that we refer throughout this book.

Those who offer a reified account of knowledge, although they may well resist such a classification, include Professor Ikujiro Nonaka of the Graduate School of International Corporate Strategy, Hitotsubashi University; Dorothy Leonard of the Harvard Business School; Varun Grover, distinguished professor at Clemson University; and Thomas Davenport. On the other side, advocates of knowledge connected to social action include professor of anthropology and director of the Institute for European Studies at Cornell University Davydd Greenwood, Professor Morten Levin of the Norwegian University of Science and Technology, and Professor Frank Blackler of the University of Lancaster. A prestigious line up by anyone's standards, but can they all be right? Recall David Snowden's paradoxical perspective of knowledge as both an object and a flow. According to this, they can. This is a debate in which no perspective can be said to be "right" or "wrong" for that matter.

The substance of debate is this: there is a world of difference between a strategy geared toward managing and leveraging knowledge as an object, lying out there to be found, a commodity to be codified, counted, stored, and even sold, and the view of knowledge as a socially constructed and distributed phenomenon. The two respective sides are fairly entrenched with, for instance, Greenwood and Levin describing the commodity view as

"petty." All of this, of course, suggests an either/or case, which according to Snowden is old thinking. But this is an incredibly complex issue. It is so complex that it is a theme to which these discussions return to throughout most chapters of this book. It is, for instance, fundamental to theories in KM, to which an entire chapter is devoted. In among all of this complexity, it is nonetheless clear that whichever agenda is adopted, it will have major ramifications on any business strategy designed to manage knowledge and on precisely what is managed.

But what of the traditional hierarchical model: data—information—knowledge? Max Boisot of the School of Administration in Barcelona suggests that data is out there in the world, knowledge is in people's heads, and information is the element that mediates between them. Can this not offer a concise understanding of the nature of knowledge, as defined in relation to data and information? Unfortunately not, but a clearer understanding of this particular model can be gained following a discussion of the issues and competing formulations of knowledge.

So far, a landscape has been imagined in which the ranks of "knowledge as object" are pitched on one side, with those of "knowledge as social action" on the other. To put these polarized viewpoints into perspective, two further questions are briefly considered: the subjective versus objective view of knowledge and the debates around organizational versus personal knowledge.

1.3.2 The Subjective versus Objective

The subjective/objective debate, which can be equated to the debate over knowledge as a "thing" versus knowledge as "action," has a lengthy tradition in the philosophical and scientific domains, with a recent trend seeing a shift away from a strict distinction between the two. A simple way of looking at this is to see the subjective as personal, open to bias and opinion, evaluation, attitude, and belief. By contrast, the objective relates to undisputable facts, divorced from bias, experience, or influence.

As background to the subjective versus objective question, consider two cases of "objective accounting" drawn from psychology: witness accounts and scientific reporting. Witness testimony is idealized as "objective," an unbiased account of what happened, and in fact, many witnesses will go to extreme lengths to frame their accounts as such. Social psychologists Abigail Locke and Derek Edwards, for example, studied President Clinton's testimony before the US Grand Jury concerning his relationship with Monica Lewinsky. According to their fascinating analysis, Clinton portrays himself as determined to focus solely on the facts of the matter—the objective, knowable account of events—and thus is able to sidestep awkward questions by appealing to his faulty memory.

Witness accounts are not just confined to courts of law: in everyday conversation, for instance, we routinely frame our reports, accounts, and descriptions of events as valid and authentic as firsthand experience. In his study of "witness accounts" by callers to a radio station, Ian Hutchby of Brunel University finds that narrators of firsthand eyewitness accounts show themselves as in possession both of the rights to recount the experience and also how it emotionally affected them, whereas the person listening to the story usually does not. If the recipient were to retell the account and couch it in their own emotional response, this would seem odd. Moreover, as Elizabeth Loftus, professor of psychology and law at the University of Washington, has shown in her empirical studies of eyewitness accounting, memory is remarkably faulty and seriously open to influence. How unbiased is that?

According to folk wisdom, scientific reporting would stand as the paragon of unbiased, objective accounting. But it is also undeniably bound to human perception. Consider an account often related in psychology texts to underline the concept of "individual differences." In 1795, Nevil Maskelyne (1732–1811), England's royal astronomer, sacked his assistant, David Kinnebrook. According to the account given by psychologists Duane and Sydney Schultz, Maskelyne noticed variances between his assistant's observations and his own of the same phenomena. He concluded that his assistant's work was inaccurate, and thus Kinnebrook became victim to one of the earliest records of human individual differences. Twenty years later, the German astronomer Friedrich Bessel (1784–1846) proved that it had indeed been a case of individual differences—the innate, unavoidable disposition of the individual to uniquely perceive the world—sufficient to result in small observational differences.

The question then is to what extent can anything ever be objective or wholly subjective for that matter? Alan Chalmers, professor of the History and Philosophy of Science at the University of Sydney, questions whether science and its knowledge can ever be considered objective. He is not alone. Michael Polanyi, referring to knowledge, argues that complete objectivity is not only a delusion, but a false ideal, claiming that "…the act of knowing includes an appraisal, and this personal coefficient, which shapes all factual knowledge, bridges in doing so the disjunction between objectivity and subjectivity" (1962: 17). What he means is that all "factual knowledge" comes drenched in personal, individual evaluation—experience and context—which renders even the most factual of knowledge as an inevitable mix of the subjective and the objective. Kenneth Gergen is even more critical of the idea that truth and knowledge are somehow lying "out there" to be discovered, citing human "individual differences" as the rug that pulls the feet from under objectivity. Everything—or mostly everything—is therefore subjective—or mostly subjective.

The objective/subjective question of the nature of knowledge takes on strategic importance in the context of knowledge as an organizational asset to be managed and leveraged. Those, for instance, who adopt a "systems approach" (i.e., with an emphasis on information technology (IT) as the primary means of managing knowledge) tend to apply an "object" label to the commodity they manage. This touches on yet another schism in KM: for some, its practice is mainly concerned with technologies and how they are used in managing knowledge, taking its roots from information management. So we see a proliferation of "knowledge repositories," "yellow pages," knowledge discussion forums, and a ramped-up increase in interest in the use of social media. The concerns of the systems approach center on effective content categorization strategies: the ability to find documented knowledge at the point of use and the people who possess the knowledge that is needed, for instance. Such technologies are valuable through their communicative potential, but can their contents really be called "knowledge"?

As might be expected, there are equally strong opposing views to the central role that some give technologies. For instance, in their analysis of case studies of information and communications technology-focused knowledge initiatives, Kenneth Grant and Umair Qureshi of Ryerson University in Toronto found high failure rates, which they correlate to initiatives' emphasis on technology. In mitigation, IT projects are notoriously prone to risk and failure but, as Grant and his colleague conclude, knowledge IT projects come with an additional set of complicating parameters and must account for the personal nature of knowledge and the importance of communities.

An early conclusion to be drawn is that the either/or approach to knowledge as object or social action, as objective or subjective, has significant consequences for what is managed or studied and their outcomes. Grant and his coworker's reference to personal knowledge leads into the next theme.

1.3.3 Organizational versus Personal Knowledge

This concerns the question of whether knowledge can only ever be personal or whether it can exist collectively, typically referred to as "organizational knowledge." Once again, polarized sides are encountered. Scholars including Michael Polanyi, and Robert Grant writing in 2002 support a personal knowledge agenda. The latter, for instance, suggests that the more productive perspective is to view collective (viz., organizational) knowledge as an aggregation of the knowledge in individuals' heads. It is precisely the relationship between the two that Grant calls into question, arguing that an emphasis on organizational knowledge simply clouds the relationship between collective and individual knowledge. This is a question considered in some depth by Haridimos Tsoukas and his colleague. In an admittedly confusing definition, they suggest that personal knowledge is the ability of the individual to operate effectively within a collective drawing on their understanding of its context and all things associated with this including its "theory," which they contrast with organizational knowledge described as a set of organizational rules that people are able to draw and act upon. What they conclude is that these are two distinct types of knowledge, implying a requirement for two equally distinctive types of management.

What these many sides of the debate tell us is that KM, in its practice, research, and theory, targets a disparate—often incompatible—array of artifacts and phenomena. Is there any commonality at all to be found? In fact, there is and it is to this that we now turn.

1.4 THE TACIT–EXPLICIT CONUNDRUM

There may be no common *definition* of the nature of knowledge in KM's literature and evidence, but there is one *structural* framework that is beyond doubt the most popular version (but short of a majority) of affairs: the tacit–explicit duality. According to the essence of the dualist perspective, explicit knowledge can be easily articulated, specified, codified, captured, stored, and generally treated as one would a sheet of paper, a computer file, a book, or a report. This sort of view of knowledge is very much behind what we have seen David Snowden describe as the first two generations of KM with their emphasis on technologies.

Tacit knowledge, by contrast, is difficult to articulate and costly to share. It is also, according to Dorothy Leonard, "sticky" in the sense of being hard to extract from its host. The implication of Leonard's viewpoint is that it is the personal, individual context that the "host" attaches to their tacit knowledge, which gives it its meaning, and further that if tacit knowledge is divorced from its context, it would lose that meaning. (Although glossed over here, the subject of context and knowledge becomes a recurring theme and particularly comes to the fore in Part Two.)

This leads into what is, arguably, the most contentious issue in KM: the conversion question. Professor Ikujiro Nonaka and his colleagues, in their formulation of the

"knowledge-creating company," perhaps unwittingly threw down the gauntlet with their theory that tacit knowledge can be accessed and leveraged through *converting* it to explicit knowledge, given the right circumstances and environment. (This particular theory is considered in detail in subsequent chapters, but, for now, the concern is with models and definitions of knowledge.) The upshot is a field of organizational practice that is broadly split between those who adopt the tacit–explicit duality, with or without embellishing elements, and those who take a different approach. There is no other single competing construct of knowledge with anywhere near the same level of influence as the knowledge duality model and its project of conversion. But where does this construct come from, and how has it achieved what Stephen Gourlay, a scholar at the UK's Kingston University, describes as such paradigmatic status, albeit not a "universal" one?

Nonaka is widely credited with introducing the tacit–explicit dualist structure of knowledge to the KM field in the early 1990s, drawing on the work of Michael Polanyi. However, it would not gain traction until the publication of a subsequent version, the *Theory of the Knowledge-Creating Firm*, in 1995 along with coworker, Hirotaka Takeuchi. Interestingly, David Snowden suggests that the reason for the delay in excitement over the new theory is explained by the then commitment to process reengineering, which was in "full flow" in the earlier part of the decade. Whatever, with the 1995 publication, the KM machine was given a fresh momentum. Perhaps it was seen as a palliative to what Snowden describes as a growing organizational disillusion with the earlier reengineering fashion and perhaps the realization that in replacing man with machine, the baby had been quite literally thrown out with the bathwater. In this and subsequent works, the theory of the knowledge-creating firm became inextricably bound to Polanyi's philosophies, which are largely focused on the nature of personal knowledge in the context of the exact sciences.

As a result of Nonaka's (and colleagues') work, which locates the tacit–explicit model as a central component, Polanyi's ideas have become among the most widely referenced in the KM field, second only to Nonaka himself. His ideas have also become what could be described as the most misrepresented and misunderstood, as will be discussed presently. In short, the fact that Nonaka bound his influential theory to this particular knowledge structure has proved significant, and in particular his emphasis on the importance of the tacit component has impacted the directions of the field—both in practice and academia. In that sense, tacit knowledge has become something of a chalice cup.

Predictably, the whole idea of a tacit–explicit structure of knowledge, especially the idea that tacit and explicit knowledge can be converted from one to the other, and then shared, has come in for a barrage of criticism. It has been pointed out, for instance, that to hold up tacit knowledge as the most valuable form on the one hand while on the other to insist that it needs to be converted to make it explicit poses something of a contradiction. Logic suggests that if it is no longer "tacit," then it is no longer of value. Researchers Ulrike Schultze and Charles Stabell, who adopt this line of reasoning, also point out that once converted tacit knowledge can be copied resulting in a risk to competitive edge. However, this is perhaps a circular argument as it first assumes that tacit as a distinctive structure of knowledge exists and second that it can be converted to explicit. Neither of these two assumptions should be taken for granted in the absence of supporting scientific evidence. But this reasoning does suggest that Nonaka's model is flawed.

A further critical theme, and one more widely reported, is the suggestion that Professor Nonaka and his colleagues, while drawing quite openly on Michael Polanyi's hypothesis for their construction of knowledge, have in fact misinterpreted and even misrepresented

the Polanyi's claims in respect of knowledge, particularly of the tacit dimension. If this proves to be the case, does it not crumble the foundations of the theory of the knowledge-creating firm? A question, though: Nonaka and his colleagues' work may be frequently criticized for "misrepresenting" Polanyi's ideas with consequences for the validity for their model, but that criticism assumes that Polanyi—in any accurate understanding—is right.

1.4.1 What Did Polanyi Really Say?

It may seem "off message" to slip so easily into a discussion of the works of a scientist and philosopher who was not even specifically addressing the concerns of KM. But, as noted earlier, KM refers to his work so frequently as to suggest that he was in fact writing with a management theme in mind, which, of course, he definitely was not. Polanyi's concern is to dispense with the objective, impersonal ideals of scientific detachment in favor of recognizing "knowing" as an art in which the skill of the knower is a fundamental part of scientific understanding. What he is proposing here, according to this interpretation, is that it is the scientist's *participation* in both discovery and validation of knowledge, which is itself a part of that knowledge and of the science. Consequently, science and its facts—knowledge—can never be entirely objective, because they will always include a personal *subjective* component.

Think of it like this. A market researcher canvasses opinions about a new product, which she does using semistructured questionnaires. She herself conducts the interviews, transcribes the results, and interprets the findings. At all of the following points of the exercise, the researcher could be shown to have, perhaps unknowingly and unwittingly, introduced her own personal interpretation of events based on past experience, beliefs, and attitudes:

- The selection of the research topic
- The selection of the interview candidates
- The selection and presentation of the questions
- Her own involvement in the interview
- The transcription of the interviews
- Interpretation and reporting of the findings

At any of these points, bias can slide in. The same could be argued about any research involving questionnaires, interviews, and even political polls, for instance. Why? For the simple reason that humans by their nature are designed to operate in a (un)conscious state of constant "sensemaking." Everything you know, everything you think you know, everything you can do, or think you can has arrived with you via your sensemaking filter—your perceptual senses. This notion of "sensemaking" is returned to in more depth in later chapters. For the present, and with this point made, what does Polanyi have to say on the subject of the "tacit"?

In the course of fashioning his claims over the nature of subjectivity and objectivity, he particularly focuses on tacit knowledge. He variously describes this as personal, practical, ineffable, indefinable (sic.), instrumental, and residing in subsidiary awareness. It is, in sum, unspecifiable. Thus, in his own words, "(A)an art which cannot be specified in detail cannot be transmitted by prescription, since no prescription exists" (1962: 53).

If that is Polanyi's brand of tacit knowledge, what of its relationship with the other kind, explicit? Recall Nonaka and colleagues' tacit–explicit model and the imperative to convert one to the other, a *distinction* that is widely subscribed to. A substantial number of those who disagree with this model, including Paul Duguid of the University of California, Berkeley, have interpreted Polanyi's claims as denying any divorce between the two phenomena: accordingly, all knowledge contains both tacit and explicit elements. In his railing against economists' reductive treatment of knowledge to the level of "widgets," Duguid persuasively reasons that the tacit cannot be reduced to the explicit. Moreover, the tacit is essential to engaging an understanding of the explicit. This he describes as evidenced in a long list of philosophical thought reaching back to the ancient Greeks several thousand years ago, and which understands what we call tacit knowledge as the underpinning cognitive dimension that mediates—"tells us how to use"—what we refer to as codified or explicit knowledge. The tacit, in this interpretation, is the key to unlocking the explicit.

Unfortunately, Polanyi does not offer an explicit, "cards on the table" explanation of the tacit–explicit relationship in a way that would have relevance and resonance for KM. Instead, a scenario that he uses to talk *around* the subject might shed some light. Imagine you learn to play the piano: at first, you concentrate on your fingers on the keys and their movement from one key to the next. Eventually, after much practice, you can play the piece without looking at your hands or even thinking about tapping keys. What began as the acquisition and application of what might be called explicit knowing—knowing that is necessarily in full focal awareness because you must concentrate on what you do and consciously apply the task's rules—has become tacit knowing. That is, knowing that has become subsidiary to conscious awareness to the extent that if you switched your full conscious attention onto your playing, you would probably lose your ability to play. ("Focal" and "subsidiary" awareness are Polanyi's terms.) So, from this scenario, it would appear that the explicit can become tacit through practice and experience, but the minute you attempt to "convert" that to explicit, you risk a reduction or loss in ability. However, I would suggest that the subtlety here is that both tacit and explicit knowing are implicated from the outset of the task: what changes is your awareness.

There is whichever way you approach it a lot of confusion over what Michael Polanyi actually meant. What is clear is that the majority of scholars and practitioners in the field approach knowledge as comprising two or more types, of which tacit and explicit are the primary classes. The main point of debate is over the nature of tacit knowledge, and whether this can be harvested and converted to the explicit, as many suggest. Or, as others claim, such attempts at "managing" the tacit are wholly inappropriate and a wasteful undertaking. The tacit question is returned to in a later chapter, where the discussions turn to a different scientific discipline for a more evidence-based understanding of the phenomenon. Now however consideration is given to some of the alternative accounts to the tacit–explicit duality.

1.4.2 The Importance of Context: What Context and Whose Context?

There is an argument that places "context" in the role of the most important aspect in any definition or understanding of knowledge. Cambridge University scholar Mark Thompson and his colleague, for instance, propose context to be an embedded knowledge element, and that this is what Michael Polanyi meant when he described the structure of knowledge

as containing a "personal coefficient." According to this formula, knowledge, or knowing, is meaningless without its context, and that context is specific and unique to the individual. But where does this context come from, and how does it emerge?

For some, context is personal, while others, for instance, Charles Despres, Daniele Chauvel, and Ganesh Bhatt, seem to be positioning context as a referent to environment or culture. From the standpoint of context as personal, which position Polanyi arguably adopts, for knowledge to have meaning, it must be filtered through the mesh of the individual's beliefs, experiences, expectations, and so forth. These may be organizationally shared meanings, or they may be unique to the individual. For instance, Mark Thompson and his colleague understand context from an organizational perspective, seeing it as comprising the shared meanings and experiences of organization members, implying an "organizational cultural context." Similarly, Davydd Greenwood and his colleague define knowledge as inherently collective ("knowing"), socially constructed, and distributed, embedded in the individual's understanding of how to act in the world, all of which can be interpreted as representing "context." Based on these accounts, it is clear that the notion of "context" is used in a broad and sometimes ambiguous sense. As something of a tangential indulgence, an interesting perspective on "context" can be drawn from psychological studies of memory.

The topic of memory has one of the longest traditions in psychology, and Alan Baddeley has made it his particular field of expertise. In his book *Essentials of Human Memory*, based on decades of research, Baddeley hypothesizes that we store our knowledge in a mental structure known as "semantic memory," and much of this is not capable of expression in words. This rings considerable bells with Michael Polanyi's perspective on tacit knowledge. For Baddeley, context is not so much the influence of the external environment, but rather our interpretation of it using the tools of our stored knowledge. For instance, referring to our interpretation of words, knowledge of the world that exceeds in most cases the boundaries of words' meaning is essential to understanding. This has correspondence with Paul Duguid's formulation of an uncodifiable substrate, which tells us how to use the "code."

Taking an example of context at work, Baddeley reports on a number of studies by Endel Tulving, whom he describes as one of the most significant contributors to our understanding of memory retrieval. In one of these, participants were presented with a word to be remembered (the target word), simultaneously associated with a "cue word," which had a loose association with the former: for example, if "city" was the word to be remembered, it might be associated with "dirty" or "village." In retrieval tests, the findings persuasively demonstrate that participants given the cue word were far more able to accurately retrieve the target words than those who were not. In this sense, the cue word acts as a kind of context. In the broader sense, context informs and influences meaning—our understanding and interpretation of our knowledge of the world. In other words, sensemaking is context driven.

In a more elaborate experiment of context-dependent memory recall, Baddeley and his colleague, D.R. Godden, got divers to learn word lists in one of two environments: underwater or on dry land. Their findings show that words were more accurately recalled when divers were in the environment in which the original learning took place and worse in the opposite environment. It was not the physical environment that was doing the influencing, but rather the participants' interpretation of it. Remove the context, and you remove the "sense." To quote a story recounted by Baddeley and his coworker, a man is found to be

an expert dancer in a small room containing a trunk, around which he must negotiate his moves. With the trunk removed, his performance became compromised.

The intention of including these couple of examples from psychology is to point to these topics of research and investigation as a source of relevance and interest for the present discussions.

So what has all this to do with a KM perspective on context? According to Kenneth Gergen, "(S)seeing is a theory-laden undertaking" (1991: 92). This understanding of context (here conceptualized as the context of observation), both in terms of cueing access to stored knowledge and as binding knowledge to experiential, cultural, environmental, and social factors, renders it unique to the individual. People may well share common knowledge within an organization, for instance, and they may even coconstruct knowledge. However, no two people will ever possess identical knowledge, but shared *elements of context* enable a platform for shared meaning. As an idea, this is not so very far apart from many of the formulations around knowledge discussed in KM.

It is worth noting that Nonaka and his colleagues feature the importance of context in their theory of the knowledge-creating firm in various guises. In 1994, Nonaka proposes a cycle of knowledge in the form of a continuous spiral in which the tacit becomes the explicit, and vice versa, which is performed on the stage of social practice. In a later work, along with coworker Noboru Konno, this idea is developed through the introduction of "Ba"—a uniquely Japanese concept of social space that can be physical, mental, or virtual—drawing on the ideas of a Japanese philosopher. The point is that in this monumentally influential theory, "social practice" can be equated to what others would call "context," but this aspect of the theory tends to be overlooked. We now turn to social action in more detail.

1.4.3 A Preference for "Knowing" as Action

Writing in the mid-1990s, the start of Snowden's second generation of KM, Professor Frank Blackler proposed a rather radical departure from the "traditional rational-cognitive" formulations of knowledge. What he argued for was that it would be more appropriate to approach knowledge as action that people do rather than something that they possess—thus "knowing" rather than knowledge, action rather than commodity. Central to this argument is the idea that if "knowing how" is located in action, then a direct link is created between action and knowledge.

Blackler argues that the traditional conception of knowledge as abstract, disembodied, and formal is unrealistic. Instead, he proposes five "images" or types of knowledge, building on work by H. Collins published 2 years earlier: embrained, embodied, encultured, embedded, and encoded. Note that Blackler's interpretation of encultured knowledge as socially constructed and open to negotiation, which is consequently heavily dependent on language, introduces a "discursive" aspect. The difficulty with Blackler's arguments stems from an apparent contradiction between suggesting a move toward viewing knowledge as action and his primary observation of a shift away from embodied and embedded knowledge toward the other knowledge "images." It is not clear whether Blackler is arguing for a reversal of this shift or not. The original work by Collins is worth a visit in its description of two types of human action—regular and behavior-specific acts: he argues that the former underlies tacit knowledge and is centered around "rule following" and "rule establishing," which are hard to describe and transfer, whereas the latter are

decontextualizable. He also implies the action orientation of discourse. This idea is increasingly developed as the present work progresses and represents the "main event" in Part Two.

As we saw earlier, a preference for "knowing" as opposed to "knowledge" finds support elsewhere, notably in the works of Davydd Greenwood and his colleague and Paul Duguid. These scholars particularly link the idea of "knowing" to tacit knowledge, which, similarly to Collins work, underlines its action orientation. Such a conceptualization of knowing has obvious implications for KM's project.

Not unexpectedly, the relationship between "know how" and tacit knowledge—if indeed there is one—has attracted its share of debate and contradiction. On one side of the debate, "know how" is explicitly mapped to tacit knowledge: they are one and the same. The opposite side denies such a relationship: for instance, Stephen Gourlay, in his criticism of Nonaka and his colleagues' conceptualization of knowledge and their knowledge creating model, proposes that "know how" and "know that," rather than equating to tacit/explicit, represent two different types of behavior—everyday and reflective. In contrast, John Seely Brown and Paul Duguid, for instance, map "know how" to tacit knowledge, ascribing it with the same properties (not easily transferable, a product of experience). Interestingly, they also note that "know how" is embedded in practice. Such a perspective is not surprising given the importance and emphasis that Brown and Duguid place on communities of practice in the development of new knowledge. The reader's attention is drawn to Gourlay, Brown, and Duguid's use of the terms "know how" and "know that," compared with, for instance, Greenwood and his colleague and Duguid's (writing several years later) preference for the term "knowing." Practically and arguably, the only difference between these lies in the inference of action in "knowing," and the potential for reification in consideration of "know how" for instance.

What we have seen thus far is a multifarious, often confusing, display of perspectives on the structure of knowledge, so much so that it is little wonder if the majority of those interested in KM adopt the, by comparison, simple and clear-cut tacit–explicit dualism. There are two more principle themes to consider in the present discussion of the nature of knowledge, before we might attempt to draw some conclusions—semantic frameworks for knowledge and the hierarchy of knowledge—mentioned earlier in this chapter.

1.5 FRAMEWORKS OF MEANING

Research suggests that most definitions of knowledge in the field of KM tend to be—to a larger or lesser extent—structural in nature. A significant number hang their clothes on a dualist model. But what of semantic definitions? Beginning with the most influential account, how do Nonaka and his colleagues semantically define knowledge?

Adopting what he describes as an essentially Western view of knowledge, Nonaka defines it as the dynamic process of justifying personal beliefs on the road to seeking truth. Thus, in his knowledge-creating model, knowledge is "justified, true belief," which was in fact originally proposed by Plato (427–347 B.C.), the Greek philosopher. However, in a departure from the traditional perspective that gives primacy to "truth" according to Nonaka, he chooses to emphasize the components of "belief" and "justified." Gourlay, a critic of Nonaka's theory as noted earlier, points out that justification can be based on false premises, in which case "justified belief" can be wrong. So, to suggest that justified

beliefs are true, and always true, when they are fallible conjures a meaningless state of affairs. Note that Nonaka, Umemoto, and Senoo subsequently criticize the "Western" definition of knowledge for being deficient, reformulating it as meaningful information that consists of a true, justified belief and/or an embodied technical skill.

Does Nonaka's definition of knowledge as "justified true belief not sound rather personal? Does it not conjure a sense of knowledge as some private, inner mental construct, hidden from casual view? If KM and knowledge managers are concerned with identifying, leveraging, storing, and applying "justified true beliefs" for the economic and competitive benefit of the firm, does it not begin to feel awkward?

More complex semantic definitions of knowledge frame it as a mix of truths, beliefs, perspectives, concepts, judgments, expectations, methodologies, and know how. Haridimos Tsoukas and his colleague, Efi Vladimirou, in their discussion of the relationship between personal and organizational knowledge, point to one definition in particular, which describes knowledge as a mix of "framed experiences," spiced with values, information from context and expert insight, and which springs from people's minds, but is often embedded in the paraphernalia of organizations—practices and routines, document, repositories, and so forth. This, they suggest, packs far too many things into a definition of knowledge, risking a specification of the phenomenon that is so wide-ranging that it renders it useless. Neither of these definitions arguably leads to a comfortable, commonsense, and *usable* perspective on knowledge. Will a consideration of hierarchy lead to a more acceptable perspective?

1.6 A HIERARCHY OF KNOWLEDGE

The hierarchical perspective of knowledge considers its relationship with and differences from data and information. Again, we find two polarized ends of a spectrum with, at one end, Stephen Gourlay, for instance, warning against any formula for knowledge that includes a reference to information. At the other, Nonaka sees no problematic distinction between information and knowledge: information is imaged as a flow of messages, whereas knowledge is created and organized by this flow, which is grounded in the commitment and beliefs of the individual. Recall David's Snowden's third-generation conceptualization of knowledge as a "thing" and a "flow."

The traditional hierarchical model has data in the lower tier, with knowledge occupying the top, and an implied flow from bottom to top. Much of the debate in this area centers on the relationship between the tiers. One question we might ask is this: can this flow ever take the reverse direction? Can knowledge transform into data? This would seem illogical, but, in theory, and by rights of this hierarchical model, one could take a chunk of knowledge and reduce it to its data elements.

The attraction of the model lies in its simplicity, and it is consistent with what might be described as a content management—or systems—approach to KM. This leads to the suggestion that the hierarchical model of knowledge, coupled with Nonaka's influential proposition of the dualist tacit–explicit structure of knowledge and the imperative to convert one to the other, is behind the perceived dominance of the technology-centric view of KM, particularly in practice. Perhaps fortunately, and as implied in the transition through three generations of KM, there has been a notable change in emphasis in recent years to a more people-centric focus.

A more sophisticated view is offered by Vincent Barabba of the General Motors Corporation and his colleagues. First off, they propose a hierarchy that includes "understanding" and "wisdom" in addition to the conventional line up, and they warn against the "huge mistake" of ignoring the distinction in meanings between them. Their claims favor a view of the organization as a system that cannot be reduced to its parts and that it is the interaction between the parts, like an activating network, that defines the organization. In their book, data must be processed into information in order to be of any use; information is contained within descriptions (who, where, when, what, how many); knowledge is contained in instructions (how) and awareness (I know who I am); understanding is contained in explanations (answers to "why" questions); and wisdom is concerned with "effectiveness," that is, the value of the outcome of behaviors. In essence, they wrap KM into organizational learning. This is a commonsense approach, one that has a great deal of synergy with the theory and perspective that is developed across the span of the present book.

1.7 SUMMARY AND CONCLUSIONS

At this point, one could be forgiven for wishing to quietly close the door on the nature of knowledge, ignoring the complexities and difficulties in the debate. Whatever it is, as argued by the psychologist Stephen Suddendorf in his account of the "gap" between humans and nonhuman species, it is in our nature to be in a continuous state of learning new knowledge and, through sensemaking, transforming that knowledge uniquely into our own. Human culture throughout known recorded history is thematically watermarked with the urge to share what we know with others. Yet its nature has been the source of debate for just as long, and with no immediate consensus—certainly within the field of KM—remotely on the horizon. It has been shown that the academic field and practice of KM is not immune to these issues. Some of the core questions center around the notion of knowledge as commodity contrasted with knowledge as an accomplishment in social interaction, its objective versus subjective nature, and whether it can be both organizational and personal.

There is, however, one structural formulation of knowledge that has attracted pervasive popularity—the tacit–explicit dualist model—originally introduced by Professor Nonaka and his colleagues. This draws on the earlier work of the philosopher-scientist, Michael Polanyi, yet many have criticized this structural formulation—and the theoretical model of knowledge creation that it underpins—for misrepresenting Polanyi's ideas. Alternative perspectives on knowledge emphasize the importance of context and a preference for "knowing how" as a form of social action. What has also been sewn is the seed of an idea of knowledge connected to discourse as action. This marks the beginning of what will emerge and develop across the following chapters.

As with the tacit–explicit dualist model, Nonaka and colleagues' semantic explanation of knowledge as "justified true belief is also widely popularized, although more complex descriptions of knowledge are also proposed. Perhaps the latter should not be so easily dismissed as being overly complex: they, at least, give consideration to the pervasiveness of knowledge—particularly the view of knowing as action—in all aspects of everyday life. To what extent are "knowing action" and "knowledge" requisite facets of every cognitive and physical action that humans do?

So what can be concluded from all of this? It is, perhaps, too early on in the investigation to arrive at any sensible, evidence-based, evaluation of the nature of knowledge, and one relevant to the organizational practice of KM. But one conclusion that can be drawn is that "knowledge," whatever it is, is important to the modern organization as it attempts to make its fortunes in what Nick Bontis describes as a sea of change and turbulent times. That being the case, an understanding of the constitution of knowledge takes on a paramount urgency. In the next two chapters, some of the other key issues and debates that infuse KM are explored, beginning with a consideration of what constitutes KM itself, finding mixed perspectives.

FURTHER READING

Collins, H. (1993). The structure of knowledge. *Social Research*, 60, (1): 95–116.

Duguid, P. (2005). The art of knowing: social and tacit dimensions of knowledge and the limits of the Community of Practice. *The Information Society*, 21: 109–118.

Nonaka, I. and Takeuchi, H. (1995). *The Knowledge-Creating Company*. New York: Oxford University Press.

Polanyi, M. (1962). *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: The University of Chicago Press.

Tsoukas, E. and Vladimirou, E. (2001). What is organizational knowledge? *Journal of Management Studies*, 38, (7): 973–993.