Phenom Cogn Sci
DOI 10.1007/s11097-011-9251-x

Tacit knowledge management

Rodrigo Ribeiro

Abstract How can we identify and estimate workers' tacit knowledge? How can we design a personnel mix aimed at improving and speeding up its transfer and development? How is it possible to implement tacit knowledge sustainable projects in remote areas? In order to answer these questions, it is necessary to distinguish between types of tacit knowledge, to establish what they allow for and to consider their sources. It is also essential to find a way of managing the tacit knowledge 'stock' and distribution within the workforce. In short, a conceptual framework is needed to manage tacit knowledge. Based on previous works and 2 years of action research, this paper introduces such a framework and describes its partial application to support the pre-operational training and hiring in a large industrial plant in Brazil. Two contributions emerge from the research. First, the concept of 'levels of similarity' is introduced as a means to qualify the experience of workers and estimate the associated tacit knowledge. Second, the capability of carrying out three types of judgement properly and speedily is put forward as being a core ability of those who possess what has been called 'collective tacit knowledge' (Collins in Organ Stud 28(2):257–262, 2007). In practical terms, the results indicate the opportunity for companies to capitalize on the experience and tacit knowledge of their workers in a systematic way and with due recognition. Ultimately, positive impacts are expected in their absorptive capacity as well as in their management and human resources systems, accident prevention, productivity and the development of sustainable projects in remote areas.

Keywords Tacit knowledge management · Knowledge transfer · Learning · Practice · Sustainable development

The methodology described in this paper was protected by the *Universidade Federal de Minas Gerais* (*UFMG*) through a Patent Application.

R. Ribeiro (✉)
Department of Production Engineering, Universidade Federal de Minas Gerais—UFMG, Avenida Presidente Antônio Carlos 6627, CEP: 30161-010 Belo Horizonte, Minas Gerais, Brazil
e-mail: rodrigoribeiro@ufmg.br

Published online: 22 January 2012

Introduction

The idea of 'tacit knowledge' has been with us for decades and particularly since Polanyi (1983 [1966], p. 4) first stated that 'we can know more than we can tell'. It refers to a kind of knowledge that human beings only develop through the experience gained over years and which cannot be passed on by a set of instructions, manuals and so forth. There is some controversy about what the quote above actually 'tells' us and what follows from it—which is itself quite an ironic fact.

One view is that tacit knowledge is opposed to 'explicit knowledge'; this is generally defined as a type of knowledge that 'can be articulated in formal language, including grammatical statements, mathematical expressions, specifications, manuals, and so forth' (Nonaka and Takeuchi 1995, p. viii). It is as if there were a 'verbalization' or 'codification' barrier which would define what is tacit and what is (or can become) explicit. Another view is that 'tacit and explicit knowledge are not two ends of a continuum but two sides of the same coin: even the most explicit kind of knowledge is underlain by tacit knowledge' (Tsoukas 2005, p. 158). This approach is endorsed by Polanyi (1969, p. 144) himself, when he says that 'explicit knowledge must rely on being tacitly understood and applied'.

Another, even more profound but related controversy concerns what is meant by 'knowledge'. There is a long-standing debate about the nature of knowledge being: something given or socially constructed, the property of individuals or collectivities, abstract or situated, universal or local, timeless or dated, amenable to codification or tacit and a commodity or a practice (e.g. Blackler 1995; Gherardi 2000). Hence, for some academics, knowledge is still 'created by individuals' and can be 'converted' between its tacit and explicit forms (Nonaka and von Krogh 2009, p. 635). For others, the discussion of 'knowing'—'as mediated, situated, provisional, pragmatic and contested'—should have long replaced the one of 'knowledge', 'with its connotations of abstraction, progress, permanency and mentalism' (Blackler 1995, pp. 1035–1040). Cook and Brown (1999, p. 381) refer to these distinct views as a battle between the 'epistemology of possession'—where knowledge is treated 'as something people possess' and there is a tendency 'to privilege the explicit over the tacit'—and the 'epistemology of practice'—where 'knowing *as* action' becomes its focus.

This paper is a contribution to the 'epistemology of practice' approach derived from the Sociology of Scientific Knowledge. Knowledge is taken here to be the property of a social group and constitutes a 'form of life' (Wittgenstein 1976 [1953], p. 226[e]).[1] Knowledge is socially agreed, modified and transmitted; it calls for 'immersion' to be developed; it has 'boundaries' that define who is in and who is out, who is a member or not, who was socialised within it and who was not. Knowledge is also rule-governed, dated and situated: perceiving, seeing and doing things in 'the same' way (Wittgenstein 1976 [1953], p. 84[e]) or 'properly' is the result

[1] This does not mean that there are no individual contributions on how to do or explain things better; it just points out that only after it is agreed by a given collectivity that whatever is being proposed is 'better' that it becomes part of what is considered knowledge. In other words, within a sociological approach, there can be no such thing as a private rule, language or discovery. As regards 'private language', see Wittgenstein (1976 [1953], pp. 88–89[e] and Section 243 in particular), Collins (1992 [1985]) extends the discussion to the impossibility of 'private scientific discoveries'.

of a social agreement—be it tacit or not—that changes according to the circumstances and with time, as Kuhn's (1996 [1962]) 'scientific paradigms' illustrate.

Within this approach, the essence of tacit knowledge is in the ability to fully participate in a form of life—what Collins (2007) named 'collective tacit knowledge'. This means, for instance, being able to stop the 'regress of rules' with regard to applying instructions and standards as expected as well as acting smoothly or improvising within a (technical) culture (Ribeiro 2007a, b). However, tacit knowledge is also used in two other senses, the first being to point to humans' physical and sensorial skills, such as riding a bicycle or the ability of medical staff to see nuances and distinct shades in X-rays where lay people only see blurs. Second, tacit knowledge also refers to 'taken-for-granted' practices that individuals unknowingly 'grasp' in the way of becoming members of a group. The best example here is when novices do something that enculturated actors would never do but that they only realise this when they (unexpectedly) see the novices breaking the 'taken-for-granted' rules.[2] I will refer to these two types of knowledge respectively, as 'somatic tacit knowledge' and 'contingent tacit knowledge'. The challenge dealt with and described here is on how to professionally handle tacit knowledge within organisations, and our case study is a greenfield industrial plant in a very remote area of Brazil.

This paper is structured as follows: To begin with, a brief description of the context in which the research took place is presented. This is followed by the introduction of the conceptual framework and the description of the fieldwork, including the empirical problems faced, the solutions given and their impact on some of the company practices. By the end, the concept of 'levels of similarity' is introduced as the answer to a central problem: If tacit knowledge is learned through qualified experience within a form of life, it should be possible to qualify the experience of workers and estimate the associated tacit knowledge. The capability of carrying out three kinds of judgement properly and speedily is also put forward as being a core ability of those who possess collective tacit knowledge. In practical terms, the research results indicate the opportunity for companies to capitalize on the experience and tacit knowledge of their workers in a systematic way and with due recognition. Ultimately, positive impacts are expected in their absorptive capacity as well as in their management and human resources systems, accident prevention, productivity and the development of sustainable projects in remote areas.

Research context and methods

> [The Company] commits itself to hire locally at least 45% of the total workforce [estimated at 1,490 people when in full operation] until the end of the third year, and 60% until the seventh year. … In order to reach these goals, the Programme

[2] For instance, once I had the rare opportunity of enabling students to interview a senior manager in a *steel* company only for them to come back to me in the next class saying that the meeting went very well indeed, except for the fact that half the group were not allowed into the company as they were wearing casual shorts!

> of Local Workforce Qualification will include a broad range of qualifying and training activities ([Brazilian Company] *Social Action Plan* 2006, p. 18).[3]

The above quotation is part of an official document sent by a major mining company to the Secretary of the Environment in a remote state in Brazil. It is a commitment to the development of the region where the company expects to produce 50,000 t of nickel per year in a large industrial plant. Further, it is one of the items required if an Operation Licence is to be issued by the local authorities.

Here, the hiring and training of local people is a guarantee of lower salaries and turnover; professionals from other areas of Brazil demand higher salaries to go to a remote area and are likely to leave when a good opportunity presents itself. This also contributes to a decrease in the expected migration which would otherwise make greater demands on public services and on what has to be provided by the company in order to hold its employees and their families in the region. Additionally, this is the best option for setting up a sustainable project while less is spent on turning a remote area into an attractive one.

The lack of a *local* qualified workforce is a problem, however. The plant under construction is in the north of Brazil, near the Amazon Jungle and 400 km from a major city. The two cities closest to the plant, C1 and C2, are quite small (with productive populations of 11,902 and 15,399 people, respectively), and their inhabitants are poorly educated. In C1, the percentage of people over 10 years old with 'less than 1 year of education or no education' is 22.4%, and with 'less than 7 years', this is 87.9%, while these numbers for C2 are 15.9% and 83.4%, respectively. Finally, the average income is low, with 69% and 63.4% of employed people in C1 and C2 receiving less than US $650.00/month.[4]

An additional problem is that the mining company has no previous experience or expertise in running metallurgical plants in Brazil. The plant is so new to them, so technically sophisticated and dangerous that the company has to hire qualified and experienced people from outside the company and from other parts of Brazil.[5] It is also anticipated that non-local people—who agree to go to this remote area—will not be willing to stay for more than 2 to 4 years after the start-up. This period is just enough to improve their *curricula* as the plant will employ the most advanced technologies worldwide.[6]

This broad context places the local plant management in a catch-22 situation. On the one hand, it is preferable, politically and economically to run the plant mostly with local people. On the other hand, it is heavily dependent on the support of non-local experienced people—at least in the beginning.

Given the novelty of the project and the region in which the plant was intended to be built, it was decided to create a Pre-Operational Group (POG) to work through some of the main steps of the nickel project. The POG was in charge of all the

[3] The bibliographic information of this document was not included in the "References" in order to maintain the anonymity of the company.

[4] Data taken from the 'Social Action Plan' (2006) produced by the Company, which in turn bases its analysis on official data obtained from the Brazilian 2000 demographic census and local authorities

[5] To bring experienced people from the other nickel plants of the Brazilian company abroad was not an option systematically pursued during the pre-operational phase of the project.

[6] A few years ago (2004–2006), a mining site from the same company—just 115 miles away—had almost 70% turnover in the first 2 years of operation.

activities for preparing the plant so that it could enter into operation, including the hiring and training of the entire team of workers.[7] The solution given by the POG for the situation is twofold: It offers pre-operational training and also mixes local (novice) people with non-local (experienced) workers. These two options gave rise to many questions.

What is the best way of turning people with no previous industrial experience and lower education into skilful operators? How is it possible to train someone to decide, for instance, when the production should continue *regardless* of a given problem or when an operational standard does *not* apply to the situation at hand? How is it possible to teach someone to assess risk reliably? How can someone be prepared to judge what counter-measures should be taken *firstly* in an accident, or if maintenance intervention is *actually* necessary? As will become clear, only experience and its associated tacit knowledge facilitate proper and speedy action in such cases, yet how can these be developed in a systematic way?

As regards the hiring of experienced people, how do we qualify and estimate workers' distinct experiences and bodies of tacit knowledge? How is it possible to plan the mix between the novice and experienced people? How long will it take for novices to be able to run the plant as well as for experienced employees to adapt to its specifics? How can the rate between the local and non-local workers safely increase over time?

The above questions pose both empirical and theoretical challenges. For the plant, a safe, smooth and speedy start-up and operation depends on the questions being correctly answered. For academia, the questions indicate the lack of a conceptual framework with which to manage tacit knowledge with regard to its development and distribution by and within the workforce. Contributing to the solutions of such challenges has been the objective of ongoing action research in this plant.[8]

The result of the ongoing research is therefore twofold; conceptually, it introduces a framework that enables scholars, managers and practitioners to think of tacit knowledge on theoretical grounds; empirically, managing tacit knowledge has been made operational. The industrial plant in which the framework was applied took over some of the proposed concepts and incorporated them into its practices, language and management systems. Concurrently, the recognition of this asset has become a tool for white and blue-collar workers to negotiate improvements in the workforce, training opportunities and employment packages.

This paper describes the research from October 2008—when it started—until September 2010. During these 2 years, I spent, on average, 1 week/month on site

[7] The initial target was to start up nickel production in January 2009. This was then postponed to the beginning of 2010 and, due to the 2009 world economic crisis, to December 2010 (first metal).

[8] The project proposed to the company, based on the researcher's Ph.D., had the general goal of applying concepts from the Sociology of Scientific Knowledge to support and enhance the transfer of tacit knowledge in industry. Since it was the first time the POG was formally established in the company, the person in charge of it was particularly concerned to receive feedback on the implemented practices in order to improve them for the intended project and any future ones. This context enabled the approval and sponsoring of the project by the company to take place more easily. The general goal was then translated into more specific research goals while maintaining its original aim. Although the project was sponsored by the Brazilian company and full support was given to fieldwork trips and safety, no constraints at all were imposed on the researcher's academic freedom (i.e. the definition of the research question and scope, the methods and the writing up of the results).

and participated in various company meetings in my hometown.[9] The fieldwork comprised more than 120 recorded interviews with workers at all levels of the hierarchy, plus a large amount of informal conversations. Some documental and quantitative analysis also took place. Thus, although only part of this material is actually used here, it is this overall experience that is the basis for what is now reported, and this informs the discussions throughout the paper.

Tacit knowledge management in theory

Within organisations, the tacit/explicit and knowledge/knowing dichotomies discussed in the "Introduction" resemble the distinct ways in which management values and reinforces standards and methods as means for workers to perform well. I will refer to these two views as the 'prescriptive' and 'practice-based' approaches to learning and doing. In the prescriptive approach, the key point so that employees can work smoothly, uniformly and safely is to write down standards and methods—based on the 'best practices'—that are to be followed and used. Once articulated, knowledge could be handled, transmitted, stored and made available through the written word, diagrams or even 'knowledge portals'. Emphasis is then placed on the formal training and the daily reinforcement of standards as well as methods, and non-compliance to these leads to sanctions.

Conversely, from the practice-based approach, smoothness, uniformity and safety are the results of workers' abilities and skills to anticipate problems and deal with everyday variability and the unforeseen. It is the workers who support standards and methods here—not the other way around. The latter are applied, adapted or even overlooked according to the workers' experience and according to the specifics of the situation. Strict compliance to them in all situations would lead to a decrease in production and quality performance and even create unsafe conditions in certain cases. From this approach, emphasis is placed on learning through participation within 'communities of practice' (Lave and Wenger 1991), and once enculturated, workers have greater autonomy to judge and apply standards and methods at their discretion and to propose new ways of working and doing. This does not mean that manuals, operational standards and so on do not have any use but, as hinted, that they are insufficient by themselves as their correct application depends on developing the knowledge and discernment of their users. This indicates the relevance of the theoretical discussion on tacit knowledge and the practical need to manage it within organisations.

In simple terms, managing tacit knowledge is defined here as *managing who is going to work with whom, doing what and for how long*. Its purpose is to optimize the opportunities for tacit knowledge transfer and development. In practice, it means managing a broad range of aspects—e.g. organisational, cultural, social, personal, linguistic, economic, gender- and power-related, political and even geographical and architectural—in order to allow the access to and contribute to the immersion of individuals within specialist groups or facilitate the initial formation and maintenance

[9] A typical trip to the location of the fieldwork (i.e. to the industrial plant) requires three commercial flights and two connections within Brazil, taking 11 h from door to door.

of such groups. This, in turn, calls for a framework that enables the identification of the distinct types of tacit knowledge and the abilities they grant to those who have them. It also demands a way of qualifying the specifics of the practice within which they were developed and can therefore be applied back successfully. This is the subject of the following sections: "Types of tacit knowledge" and "Levels of similarity".

Types of tacit knowledge

Three main types of tacit knowledge were distinguished according to their nature and amenability to codification (Ribeiro 2007a, p. 73): somatic tacit knowledge, contingent tacit knowledge and 'collective tacit knowledge' (Collins 2007).[10] A feature of the proposed typology is the separation of what is tacit due to 'the functioning of the human body and brain' (i.e. somatic tacit knowledge) (Ribeiro 2007a, p. 71), what is tacit due to 'the historical development of the technical domain or matters of contingency' (i.e. contingent tacit knowledge) (Ribeiro 2007a, p. 72) and what is tacit due to 'its location in the social collectivity' (Collins 2007, p. 257) (i.e. collective tacit knowledge).[11]

Somatic tacit knowledge refers to the corporeal background that supports or enables a person to perform the physical counterpart of actions' intentions and to interact with the physical world; particular types of somatic tacit knowledge which may or may not be amenable to codification have been proposed (Collins 2007, 2010; Ribeiro 2007a). As far as human beings are concerned, however, the only way to develop any of these types of somatic tacit knowledge is through physical engagement with the subject matter. Thus, somatic tacit knowledge will be treated as a sole category here.

Contingent tacit knowledge stands for types of tacit knowledge that are 'embedded in the practices of a form of life but are in principle amenable to codification' (Ribeiro 2007a, p. 66). This type of tacit knowledge consists of a collection of cases which may be distinguished from each other according to their level of tacitness or, in other words, according to how aware the members of the collectivity are about having them. Extreme examples are when people are totally unaware of their knowledge but are known for doing something that works (this being 'unrecognised tacit knowledge', Collins 2001) or when people are consciously aware but want to keep 'the trick of the trade' to themselves (e.g. 'concealed knowledge', Collins 2001).

This paper does not comprise a detailed analysis of contingent and somatic tacit knowledge due to its focus on the *pre*-operational phase of the industrial project, where the actual practices of producing nickel were still not in place. However, these types of tacit knowledge are explained succinctly in this section for the purpose of the completion of the conceptual framework. The emphasis here therefore relies on

[10] Amenability to codification is here understood as amenability to reification: the possibility of creating products that can be used by enculturated actors (e.g. books and standards) or incorporated in technical devices (e.g. gears and software) that can then produce an output that is accepted within a given collectivity.

[11] When this typology was first introduced (Ribeiro 2007a), I used 'transient' instead of 'contingent' tacit knowledge but changed to the latter because it is more self-explanatory. I owe the change of terms to Harry Collins, although the definition was kept the same as 'transient tacit knowledge'. The development of this typology was possible due to more than 30 years of research on this topic by Harry Collins.

expanding the theoretical description of collective tacit knowledge and the abilities it supports as well as empirically confirming its presence and practical outcomes within the pre-operational phase.

Collective tacit knowledge enables an individual to perform actions that call for an understanding of the social context in order for them to be properly performed (Collins 2007). This knowledge is located in society and is not amenable to codification—at least not until machines are built that are able to socialise with human beings. Since we are accepting knowledge as being socially constituted and the propriety of a collectivity, the development of collective tacit knowledge is intertwined with becoming a member of the given form of life. In this sense, the work of Wittgenstein enables us to identify particular collective-tacit-knowledge-laden abilities and expand our understanding of this type of knowledge.[12]

The first ability of a fully enculturated person is that they are able to 'follow a rule' in the Wittgensteinian sense. This stands for the ability to act smoothly and speedily within a form of life, including being able to improvise when facing new and unexpected situations, to evaluate when it is acceptable to break a rule and how to do it (e.g. Leith 1986). Clearly, for social actors to be able to 'follow a rule', the rules must first be established; the establishment and alteration of rules are intrinsic to the formation of social groups and the changes within them. Hence, only full-blown members of such groups are able to modify or propose new rules (either explicitly or tacitly)—as is the case even within science (e.g. Kuhn 1996 [1962]; Collins 1992 [1985], 1989).

There is a discussion about whether or not the 'rules' within 'rule following' are tacit or can be explicated, but this does not affect our analysis. Even if the rules could all be made explicit at once, they would immediately become dated and subject to the 'regress of rules'—i.e. the fact that 'the rules do not contain the rules for their own application' (Wittgenstein 1976 [1953]). In simple terms, this means that *no rule is absolute in its interpretation and/or application*. This point was magnificently illustrated when I intentionally probed an experienced supervisor about the possibility of including *every possible case* in the operational standard of a machine (the engineers' dream):

> Supervisor: The OS [operational standard] tells you the routine, what happens in the day-by-day. You cannot have in the OS [something like] 'if by any chance the spectrometer breaks [then…], if by any chance you cannot measure the temperature [then…], if by any chance there is a problem in the pneumatic belt [then…] …
> Ribeiro: What would happen if these cases were included in the OS?
> Supervisor: You could include [them], but would the worker on the other side be able to interpret [it]?
> Ribeiro: … But now it is *written down* in the OS: '*if* by any chance this happens, do this. *If*… [something else happens do that]' and then what may happen [and should be done] would be defined for each case, wouldn't it?

[12] The interpretation of Wittgenstein adopted here is based on Winch's (1990 [1958]) work, which is broadly accepted within the field of the Sociology of Scientific Knowledge.

> Supervisor: Yes… (long pause) it may, but what if [the worker] does not know how to use these *ifs*?

The regress of rules implies that every interpretation and/or application of a rule calls for human judgement—hence, for collective tacit knowledge—at some point in the course of the action. This is true for tacit rules, for written rules contained in manuals and standards or even when dealing with fixed rules incorporated into machines. Thus, two of the activities of control room operators in automated plants are to supervise the performance of the machines and guarantee that the physical conditions and technical systems for which they were designed to work are met and maintained. In other words, when unambiguous and fixed rules are incorporated into machines, the judgemental part of their application is transferred to other humans, such as their designers, users and maintenance workers (Ribeiro and Collins 2007).

To follow a rule is therefore intertwined with the ability to make 'correct judgements' (Wittgenstein 1976 [1953], p. 227^{e}) within a form of life. This is because 'the notion of following a rule is logically inseparable from the notion of making a mistake … [However,] a mistake is a contravention of what is established as correct' (Winch 1990 [1958], p. 32) and what is 'correct' is always a matter of ever-changing social conventions. Thus, Wittgensteinian rules are here understood not as 'propositions' but as 'social institutions or social customs or social conventions' (Bloor 2002 [1997], p. 5).

It follows from the above that the presence of collective tacit knowledge can be spotted in daily practices if one looks for any kind of 'judgement' being made by enculturated actors. There are three kinds of judgements which only members of a given form of life are able to perform properly, bearing in mind their lived understanding of the (technical) culture within which the yardsticks were established. They are the 'judgement of similarity/difference', the 'judgement of relevance/irrelevance' and the 'judgement of risk and opportunity'.[13]

The 'judgement of similarity/difference' underlies the ability to identify *what is considered* '*the same*' *as well as violations of tolerance* (e.g. mistakes, improprieties and problems) in rule following situations and outcomes. In industry, this means being able to judge, for instance, if something (e.g. a variation, a tendency or a given datum) is or is not a problem in the first place, if a given problem is or is not 'similar to one whose solution path is known … [and] when to apply a particular [or an alternative] procedure' (Gorman 2002, p. 228). This type of judgement also underlies the ability to *create contrast* (similarities versus specifics) between situations, scenarios or technical proposals, to *provide reliable estimations* based on past experiences and to *make correct* '*approximations*' (Kuhn 1996 [1962]) *in a field*.

[13] A rhetorical question posed by Wittgenstein highlights that even what is considered 'the same' depends on what is *agreed* as being 'the same': "But isn't the same at least the same? We seem to have an infallible paradigm … I feel like saying: 'Here at any rate there can't be a variety of interpretations'" (Wittgenstein 1976 [1953], p. 84^{e}). The same line of reasoning applies to the other notions under assessment—i.e. 'relevance', 'opportunity' and 'risk' (with possible exceptions to physical risk). See "The turning point: the operational readiness workshop" and "Making correct judgements and the two approaches to learning and doing" sections for an empirical confirmation of this point. See Dreyfus (1979 [1972], 2009) for a discussion of 'relevance' within expert systems and the Internet and Collins (1992 [1985]) for a discussion of judgements of similarity and difference in the replication of scientific experiments.

The 'judgement of relevance/irrelevance' is the ability to *locate and attribute value* to events, claims, artefacts and people within the current and past history of a given form of life.[14] This judgement enables enculturated actors to *prioritise correctly*, to '*retrieve selectively*', to *evaluate who is who*—and *who to trust*—to *identify key changes/tendencies* and to *weigh the pros and cons* between options.[15]

In some cases, judging relevance/irrelevance presupposes or encompasses the 'judgement of risk and opportunity', i.e. the ability to *evaluate the* (short-, medium- or long-term) *consequences* of ongoing or future actions or events within a form of life. In industry, this enables experienced professionals to *anticipate problems and accidents*, to *decide the right* '*timing*' (e.g. 'moving forward' or stepping back), to *assess risk*, to *identify promising opportunities*, to *analyse trade-offs effectively*, to *make the best compromises*, to *frame a problem properly* and to *choose the best solution* for it.

Speed in acting correctly is the outcome of developing the above types of tacit knowledge. The higher the expertise, the faster humans act and react to situations (Dreyfus and Dreyfus 1988). Thus, a quick risk assessment of a situation may be the difference between having a potential accident or a real one in the same way that a quick assessment of an opportunity may be the difference between being the first or the second to innovate or to spot a new market. Certainly, speed in acting correctly has a positive impact on the efficiency and efficacy of any organisation—whether this concerns the use of physical abilities (i.e. somatic tacit knowledge), decision making (i.e. collective tacit knowledge) or the selective retrieving of lived experiences and cases (i.e. contingent/collective tacit knowledge).

Levels of similarity

Experienced people are therefore the best guarantee of a successful and safe start-up: They already have all the above types of tacit knowledge intertwined and incorporated into their practices and bodies. During the fieldwork, however, problems emerged with the challenge of estimating the amount of tacit knowledge within the existing workforce. In order to achieve this, tacit knowledge estimation should take into account the past experiences of the experienced workers, but only those experiences connected to their future tasks in the new plant. In other words, the question is: *How is it possible to qualify experience?*

It is not possible for someone to 'just have' tacit knowledge or have tacit knowledge *in everything*. It is only possible to have tacit knowledge *in something*—and this 'something' must be qualified. For instance, judgements of similarity/difference, relevance/irrelevance and risk/opportunity are the *same* in any area with regard to their *nature*, but their *outcomes vary* from one field to another as well as from one task to another. Therefore, it is not enough to say, for instance, that someone has 20 years of experience in maintenance; there are as many 'experiences in

[14] When applied to values, the judgement of relevance/irrelevance becomes a 'moral judgement'. Although this type of judgement is central for shaping established values and practices, its discussion is beyond the scope of this paper.

[15] The task of 'filtering the data' (Collins 1989) is an example of 'selective retrieving'. To know what counts and does not count as the right data calls for an understanding of how it is valued within a given area or for ongoing discussion as well as knowing who has produced it.

maintenance' as there are areas of maintenance (electrical and mechanical maintenance, civil engineering, automation and instrumentation), types of maintenance (corrective, preventive and predictive maintenance), roles in maintenance (e.g. inspector, electrician, mechanic, vulcanizer, maintenance planner and so forth), types of equipment and different brands (which matter sometimes), varying equipment capacities and complexities (e.g. level of automation), different ways of organising maintenance (departmental or central maintenance) and so on. Hence, it is impossible to qualify experience based on every single detail of one's work—where this so, there would be one single 'experience' for every individual! The way out of this conundrum was to adopt a sociological approach within which only the most significant specifics in each of the existing functions within the plant and the experience that could lead to their development would be taken into account.

The concept of 'levels of similarity' has been proposed to qualify the experience of experienced workers. The idea is to compare how similar a person's previous working experience and workplace are to the function the person is expected to perform: The higher the similarity, the closer it is to the experience needed and vice versa. The term 'similar' is central here for it embodies the idea of 'following a rule': What is or is not similar calls for a judgement that is dated and varies according to the field. Hence, judgements of similarity/difference and relevance/irrelevance will be more or less precise depending on how enculturated the person judging is, as will be illustrated in the next section.

Three levels of similarity—high, medium and low—were initially defined. Individuals rated as being of high similarity were those who have worked in nickel plants using similar *processes* and *equipment* to the ones adopted in the plant under examination. People with this experience were assumed to possess the highest amount of tacit knowledge; they would, however, need to adapt to some specifics of the new plant (e.g. the high level of automation), but most of what they knew could be applied in principle.[16] Medium-similarity individuals, in turn, have worked in industrial plants which use *similar equipment* to that found in the plant, but *within a different process*—or vice versa. The best example of this is those who have worked with a kiln in cement plants for a long period.[17] These people possess tacit knowledge about the equipment operation and maintenance, but they need to learn how its use *within* the nickel process affects the operational parameters they were used to. Conversely, low-similarity employees (LSEs) have worked in industrial plants which do not use similar equipment and processes to the ones adopted (e.g. the food industry). In this case, it is expected that the person will have only limited notions of safety in industrial settings, but these may be of use in avoiding simple accidents.

With levels of similarity, it is therefore possible to assess whether or not the tacit knowledge the workers had the *chance to develop* in previous jobs is more or less similar to the tacit knowledge they need in order to perform their future role at the new plant. Of course, whether or not those workers considered of high similarity *do* have the

[16] I do not mind saying that someone 'possess' tacit knowledge when it has the meaning that it is the result of becoming a representative member of a form of life. If the person stops participating in the given form of life, what she or he 'knows' may become obsolete due to changes in the established practices.

[17] This kiln is a huge piece of equipment—in our case, a 6-m-diameter, 135-m-long refractory-lined rotating vessel or rotatory vessel—used for calcining different types of raw material.

necessary tacit knowledge can only be assessed in principle after the start-up. However, it is possible to ascertain which of the experienced workers *do not* have the necessary tacit knowledge by checking the forms of life they *did not* have the opportunity to be enculturated within (i.e. low similarity). It is also possible to ascertain which workers have *some* relevant tacit knowledge by checking to see if the previous form of life they were members of overlaps with the future one in core aspects (i.e. medium similarity).

Tacit knowledge management in practice

Mapping tacit knowledge

The nickel plant under investigation can be divided into three main areas: (a) mining, (b) dryer and calcining and (c) smelting and refining. The focus of the research concerns the Operations and Maintenance Teams of the latter two. These areas are ideal for studying tacit knowledge because this is where the most dangerous activities will take place, and they constitute areas where the company has no experience in Brazil to date. Hence, it is possible that the 'taken-for-granted' practices embedded in tacit knowledge will appear as 'delays', 'mistakes', 'bad practices' and even 'inevitable accidents'; however, this can only be identified as such *afterwards*, which frequently happens when learning occurs.

Clearly, the 'action' part of the action research aims precisely at minimising these problems by maximising the development of tacit knowledge and its presence within the workforce. Thus, the first task in the research was to identify the *sources* of tacit knowledge within the project. Two of these were analysed at the beginning of the fieldwork: the pre-operational training of novice employees (hired locally) and the presence of experienced people (hired in other parts of Brazil). Here, the discussion is focused on the latter.[18]

The experience of the experienced workers

The experience of the experienced workers was estimated according to 'levels of similarity'. At this point, 90.3% of the working curricula with the experienced people hired for the Operations Team and 89.4% of those from the Maintenance Team were analysed. The many experiences they had undergone in their professional lives were first classified according to low, medium and high similarity, and then the periods corresponding to each of them individually were summed up separately. In short, the unit for estimating the existing 'stock' of tacit knowledge among the workforce was the total time of working experience per level of similarity. Figure 1 demonstrates the results.[19]

[18] The analysis of the pre-operational training of novices with regard to tacit knowledge development can be found in 'Levels of Immersion, Tacit Knowledge and Expertise', an essay of mine under review by *Phenomenology and the Cognitive Sciences* at the time the present article was accepted.

[19] The data for Fig. 1 refer to the professional experiences the experienced employees had before starting to work on the project and this information was also collected in January 2009.

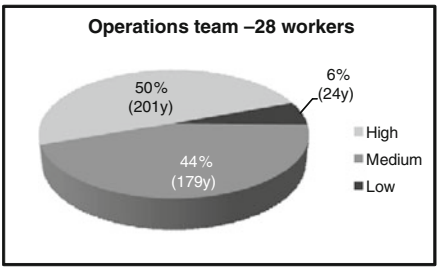

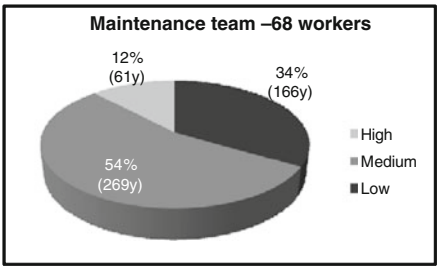

Total years of experience: 404 years

High : 201 years–Average: 13 years

Medium : 179 years –Average: 9 years

Low : 24 years–Average: 4 years

Total years of experience: 496 years

High : 61 years–Average: 4 years

Medium : 269 years –Average: 7 years

Low : 166 years–Average: 4 years

Fig. 1 Aggregate tacit knowledge of experienced workers

The previous experience of the 28 experienced workers of the Operations Team amounted to 404 years, out of which 201 years were of high similarity (50%), 179 years of medium similarity (44%) and 24 years of low similarity (6%). The average time of experience in each level was 13, 9 and 4 years, respectively. For their part, the experience of the Maintenance Team totalled 496 years for 68 workers, but only 61 years of experience could be regarded as being of high similarity (12%). The remaining time was divided into 269 years of medium similarity (54%) and 166 years of low similarity (34%). The averages were also lower compared to those of the Operations Team and would have been even lower if the few people with higher similarity experiences had not pulled them higher.

The recurring question of those to whom the above data were presented—especially from the engineers—was whether that amount of years for high- and medium-similarity experiences could be deemed sufficient or insufficient for a safe, smooth and speedy start-up. The standard answer—'I do not know'—was somewhat shocking to them, but there were some good reasons for not knowing. To begin with, it was the first time this concept was being applied and an attempt to estimate tacit knowledge was being made; currently, there is no yardstick with which to compare the numbers obtained.[20] Secondly, aggregate numbers at higher levels may hinder problems in specific areas and not provide the necessary information with which to manage tacit knowledge properly.[21] Finally, as we will see, the idea of levels of similarity was adopted and used by company employees in a variety of practical situations, allowing for far more concrete ways of managing tacit knowledge than the numbers would permit in the future.

[20] It may be the case that the 'amount of experience according to similarity' will become a good numerical indicator in the future if this type of analysis becomes a standard, so allowing for a comparison between cases within similar industries or even between different ones. Thus, a database is being built with which we will try (in the long term) to collect data from various start-ups and link them to variables that give an idea of how complex each of the participating projects are. Although this analysis will be pursued, numbers alone are not the main element for managing tacit knowledge.

[21] For instance, although the Operations Team altogether presented 'good numbers', there was one area (Utilities) that did not possess a single person of high similarity in its team at the time the analysis was carried out. This showed that the numerical analysis could lead to better results when the graphs were produced for each of the smallest 'working groups' within the Operation and Maintenance teams.

The turning point: the Operational Readiness Workshop

A 2-day Operational Readiness Workshop was held at the construction site in August 2009 to plan the hiring and training of the remaining workers for the start-up at the end of 2010. The director in charge of the project, the three general managers and all the area managers were invited to and participated in the workshop. Some highly experienced supervisors, the engineers supporting each of the operational areas and three external consultants also took part in the meeting.

The workshop began with a brief talk by the POG general manager and my counterpart at the company about its goals. A 3-h lecture followed in which I introduced the above theoretical framework in simplified terms, with accounts and practical cases from fieldwork to illustrate the main concepts. Figure 1 on the Operations and Maintenance Teams was presented as an example of how this type of analysis could help to plan and follow up the hiring and on-the-job training of the remaining employees. The participants were then grouped according to area and asked to complete two main tasks:

(a) To write a standard of 'levels of similarity' for each function of their team. This should comprise the details of the working experience for each of the three levels—low, medium and high—plus an estimation of the minimum amount of time necessary for a person to master what was specified to each level
(b) To plan the 'mix of similarity' of their teams by (1) classifying those already hired in their team as being of low, medium or high similarity on the basis of the above standard and (2) to define the desired level of similarity for each working place to be filled (i.e. of those still to be hired before the start-up)

The presence of a very experienced person who knew about the functions that would be dealt with was stressed as a requisite for the above tasks to be carried out properly. All the participants worked hard over the 2 days, and the results soon started to appear, both in the workshop and subsequently.[22]

Experts in action: defining standards of levels of similarity

In the hands of the experts, the concept of levels of similarity has been shown to allow for a refined and simple way of discriminating between experiences. This is exemplified by its application to four cases within distinct areas of the plant and can be seen in Table 1, starting at the top of the table.

Most people would say that 'an operator of a hydraulic excavator' is 'an operator of a hydraulic excavator' in any situation. Those with *laterite nickel* mining experience do not agree; in their view, a high-similarity worker for this function must have worked *within open pit laterite nickel mines* for at least 5 years. Laterite nickel mines are part of what is termed 'selective mining', that is, mines where it is possible to distinguish between the ore veins within the block being mined. Only people with 'trained eyes' are able to make such a distinction. Moreover, to work in open pit mines requires a great deal of experience due to the heavy equipment that are used

[22] With hindsight, it is possible to say that the workshop was the moment when the idea of tacit knowledge management became operational within the nickel project.

Table 1 Standards of levels of similarity for four functions

Functions	Definitions					
	Low similarity	Min (years)	Medium similarity	Min (years)	High similarity	Min (years)
Hydraulic excavator operator	Experience as an operator of a **small** hydraulic excavator	3	Experience as an operator of a hydraulic excavator in **open pit mines**	3	Experience as an operator of a hydraulic excavator in **open pit laterite nickel mines**	5
Motor grader operator	Experience as a **wheel loader operator**	3	Experience as a motor grader operator in **roads**	3	Experience as a motor grader operator in **open pit mines**	3
Control room operator calcining	Experience as a control room operator in **any industry**	3	Experience as a control room operator of **rotatory furnaces or dryers** (cement, clay, lime)	5	Experience as a control room operator of **rotatory furnaces** in the **nickel industry**	3
Turner	Maintenance and cleaning turner of pieces and tools	1	Manufacturing turner of **simple pieces and tools** (e.g. screws; axis) **from drawings and/or broken pieces/tools**	3	Manufacturing turner of **complex pieces and tools** (e.g. Crankshaft) **from drawings and/or broken pieces/tools**	5

and the associated traffic and safety issues. In sum, the experts consider it takes around 5 years for a normal person to *operate* a hydraulic excavator and to *visually distinguish* the nickel ore from the rest with success—any parallel with somatic and collective tacit knowledge is not a coincidence.

An outsider might consider that in order to be a high-similarity motor grader operator (second line, Table 1) in the same mine of the hydraulic excavator, operator would also require experience *within open pit laterite nickel mines*, but this is not the case. The experience of operating a motor grade in *any* open pit mine is sufficient to be classified as of high similarity; this function does not depend on the type of ore in the mine where they will be working but only on the type of pit.

In the case of the control room operator of the calcining line (third line, Table 1), a couple of experts argued that operating a rotatory furnace in the cement industry is actually much harder than in the nickel industry. Hence, a normal person would need at least 5 years' experience to be considered a professional operator for the former while 3 years would suffice for the latter.

The *amount of experience* is therefore *neither* connected with *nor* a guarantee of higher levels of similarity. This goes against the commonsensical idea (which we had to refute continuously when explaining the concept to those writing the standards) that lower levels of similarity would always mean a shorter period of experience or less experienced people. It also contradicts the idea that many years of experience should mean higher levels of similarity, there being operators with more than 20 years of experience who were considered of medium similarity. Finally, given the complexity of the activity, there were cases where the minimum period of experience required for medium similarity was lower than the one required for high similarity (fourth line, Table 1), but this was not the rule. In most of the cases, the experts defined the same amount of time for the two higher levels of similarity.

Finding the right mix of similarity

The idea of planning the mix of similarity was soon undertaken by the workshop participants. Based on the standards, the managers then made a rough evaluation of their current team, classifying the employees under their supervision as being of low, medium or high similarity.[23] The next step was to define how many of those still to be hired should be of high or medium similarity in order to produce the right 'mix of similarity' within the team. A team with only high-similarity people would be the optimal solution, but this was not feasible, mainly because of the high cost of hiring them and, more importantly, the scarcity of high-similarity workers in the market (for some functions, there is only one nickel industry in Brazil similar to the one being built). The managers and supervisors were then advised by the top management to plan the best possible mix of similarity for their team that would make them feel 'comfortable enough' to start up the plant. The participants adopted at least four rationales in order to reach a good mix of similarity for their teams:

> Mix by *working group* (i.e. the smallest group of people who work together towards a shared task): One supervisor stated he was going to have about 16

[23] It was a 'rough evaluation' because the managers did not have the detailed curricula of their subordinates. There were also areas where the manager and the person supporting him were not themselves of high similarity.

people working in four shifts of four people each. Thus, at least one of the four workers in each shift had to be of high similarity.

Mix by *complementarity*: This rationale was adopted by two supervisors in area maintenance, one with previous experience in mechanical maintenance and the other with electrical maintenance. Both knew they would be working in different teams, but each team had two vacancies, one for a mechanic and the other for an electrician. Their decision was to hire a high-similarity mechanic and a low- or medium-similarity electrician for the team whose supervisor had previous experience in electrical maintenance (in this way he could learn with the former and train the latter) and vice versa for the other supervisor. More interestingly, they agreed that the selection of the high-similarity mechanic was not going to be made by his or her future supervisor, but by the supervisor who had previous experience with mechanical maintenance—and vice versa with regard to the selection of the electrician (this shows the supervisors' awareness of the limits of their experience and reinforces how judgements are restricted by the experiences individuals have undergone).

Mix by *operational necessity*: Here the supervisor in charge of the area maintenance of the furnace and the refining argued that he needed two high-similarity instrumentation technicians in his team for the moments when problems in these areas occur simultaneously. There is no stock of material between the furnace and refining in process which calls for fast intervention.

Mix by *precaution*: For this, a supervisor stated that he was going to try to obtain two high-similarity workers because of future turnover common in the function he was talking about.

To find and work towards a better mix of similarity became then a method for planning and managing the composition of the teams with the clear goal of increasing both the amount of tacit knowledge of the workforce and the chances of tacit knowledge development within the working groups.

The explicative power of 'levels of similarity'

The concept of levels of similarity has proved to be relational and situated, but only when used by high-similarity workers. Their standards incorporate precise judgements of similarity/difference and relevance/irrelevance within and between the descriptions of the three levels of similarity. During the workshop, this could only be achieved by representatives of the highest level of similarity or (in simple cases) by those with a wide experience in the different types of industries concerned. Therefore, low-similarity participants in the workshop had to count on the support of their high-similarity peers or supervisors in order to write the standards. Those who tried to write them without the necessary support came up with very abstract descriptions of the levels and were not able to explain quickly and simply exactly why they had included or suppressed the qualifications present in the descriptions. Conversely, high-similarity workers, having understood the concept, wrote their standards 'naturally' and quickly and had no problem in defending them.

Given that the curricula analysis for obtaining the years of high and medium similarity of the experienced workers (Fig. 1) was not based on the more refined

standards emerging from the workshop, it can be assumed that the corresponding figures are overestimated. The criteria initially used to classify the curricula had a broader area of tolerance as to what should be considered 'high similarity' than the one defined by the experts. In addition, those who conducted this—the researcher with the help of two undergraduate students—could not properly judge the similarity/difference of what was written in the curricula compared with the adopted criteria. We are now trying to redo the calculations based on the more refined standards in order to reach the right numbers and to have an idea of the magnitude of the judging 'errors'. This has no effect on the research because this quantitative analysis only helped to highlight the differences in amount of higher similarity experience between the Operations and Maintenance Teams, which would remain anyway; however, the identification of such 'error' illustrates the analysts' inability of judging similarity/difference and relevance/irrelevance due to lack of experience - and collective tacit knowledge - in the field.

In sum, collective tacit knowledge in a given field is necessary in order to establish the standards of levels of similarity for that field and to apply them properly. This corroborates the theoretical point on the abilities a person with collective tacit knowledge possesses. On the other hand, it also demonstrates that these abilities only have a practical value within a *circumscribed* domain where a person had the chance to be immersed for an extended period of time. We now move to what happened after the Operational Readiness Workshop.

Post-workshop: social and management impacts and developments

The concept of *levels of similarity* became part of the language and practice within the plant. In fact, shortly after the workshop, managers and supervisors started referring to themselves, to each other and to other workers according to their previous experiences. Thus, it was possible to hear employees identifying themselves and others with expressions such as 'he is high similarity', 'you are medium similarity' or 'I am low similarity' with increased frequency during the following fieldwork visits. Moreover, the concept either created new practices or was a means of defending or improving existing ones. Instances of such cases can be understood as ways of managing tacit knowledge on a daily basis.

Negotiating, selecting and hiring high-similarity employees

Room for negotiation: more experts and the improvement of vacant post-conditions Planning the right mix of similarity has made some room for negotiation. For instance, there were moments where managers and supervisors realised that they were of low or medium similarity and needed someone of high similarity to support them. Thus, an area manager approached me one day in order to 'share the good news': After the workshop, he had convinced his general manager that he needed two specialist engineers to support him in areas where he had no previous experience.

There were also cases where the overall post-conditions of some vacancies—especially the salary—were not sufficient to attract high-/medium-similarity professionals for the vacant posts. The supervisors then talked to their area managers in order to upgrade the post levels, changing them, for example, from

a technician I status to a technician III. This resulted in improvements in the post-conditions and facilitated the hiring of higher similarity workers.

Change in the employment package and rules More importantly, planning the right mix highlighted the need for high-similarity workers in all levels of the hierarchical structure, but some of the vacancies at the lower levels—such as those for electricians and technicians—did not offer housing and other advantages, which would only be available from a certain level on. Furthermore, because of the high demand for housing in the region, rent for housing became higher than the salary of some posts at the lower levels of the structure!

Clearly, the supervisors were already aware of this problem before the workshop, but now they had a language and concepts—also shared by their superiors—with which to address it and back up their claims. The solution to the problem required the support of the top management of the nickel project so that the need to attract key people to low-hierarchy posts could be discussed with the Brazilian company's Corporate Human Resources team. The result was the establishment of an exception within the company: If it was proved that the worker below a certain hierarchical level needed to be of high similarity, the company would provide a house or the rent for married professionals and a room in its lodgings for single people.

Another rule that was made flexible for specific cases was to waive the company requirement that only people who had finished high school could be hired. Thus, a supervisor recalled a mechanic who had worked in the plant construction and who could be employed as someone with medium similarity (given his knowledge of the equipment) to work with maintenance. After defending this person's higher similarity, he was allowed to hire him, although the mechanic had not completed year 8.[24]

Open use of personal–professional networks Of course, the best way to find and invite high-similarity professionals for the selection process was to use the personal–professional networks of the high-similarity workers already hired. This led to a concerted action between some of the employees who came from high- and medium-similarity companies and the Human Resources (HR) team. The open use of such networks based on previous shared working experience and trusting relationships therefore became recognised and accepted within the project.[25] However, this did not necessarily mean that those passing on their contacts to the HR team were merely trying to bring in their friends and acquaintances. In other words, there was a recognition of the collective tacit knowledge of those with high similarity who knew 'who was who' and 'who to trust' within the nickel industry and other related industries.

[24] Nevertheless, those who are hired in this way must commit themselves to finishing high school within an agreed period.

[25] This strategy has led to a few complaints from other companies which were left with no other option than to promote key people within their workforce so as not to lose them. The selection through the use of personal–professional networks seems to be a way of having very cohesive teams composed of people who are trusted and committed to each other. On the other hand, it creates 'pockets of power' based on loyalty which may or may not be good for the project in the long run depending on how loyalty (or subservience) may overrule the technical aspects at stake in each case. The problem is that 'loyalty' also happens among low-similarity people.

Improving the HR selection process For those cases where no network was available or where there was a problem contacting the person in the other firm, the standards of levels of similarity were used to improve the description of the post requirements for the HR selection team. Here, managers and supervisors started to add the description and the minimum amount of experience stated in the standards in the 'Required Experience' section of the company's form that begins the whole selection process. Thus, instead of writing '5 years of experience', as would be required when hiring an operator of a hydraulic excavator, they would write: '5 years of experience as an operator of a hydraulic excavator *in laterite nickel mines*'.

The use of an appropriate level of similarity description as part of the requirements for filling a vacant post—in the light of the desired mix—goes against the idea of a generic 'operator of a hydraulic excavator' or 'mechanics' who can perform well *everywhere*. Instead, in-between qualifications leading to the selection of a more efficient professional for the case at hand could be designed and used as a guide. This indicates an opportunity for developing a fast, more technically refined and even cheaper HR selection and hiring system. It would be a mistake, however, to think that HR professionals would be able to write the description or even have the appropriate connections to find the right people. Rather, as already illustrated, the HR team and high-similarity professionals would have to work together.

The compensation strategy: hiring high-similarity support

Managers and supervisors have not always succeeded in hiring the high-/medium-similarity workers they desired despite all the effort mentioned above.[26] The main structural barriers faced have been the location of the project, with poor overall infrastructure and the high demand for skilled workers in the Brazilian labour market. Internal to the company, there were cases where (a) the salary and advantages offered were not accepted by the candidate, (b) very experienced high-similarity professionals gave up their candidacy for a post in the project due to the length of the selection process (the period within which he/she was offered a promotion or was hired by another company), (c) the candidate failed in some of the HR selection phases prior to the final interview, (d) managers and supervisors could not persuade their superiors to increase the salary offer or to include the candidate as an exception within the new rules, (e) preference was given to hire the partner of an employee in order to retain the couple in the region and (f) low-similarity managers and supervisors in charge of the selection could not properly assess the candidates' CVs in terms of their similarity. Hence, when the right mix of similarity is not achieved by hiring the appropriate people at the start, it will have to be achieved by other means in the subsequent phases of the project. One of these is to contract high-similarity

[26] Given that the implementation of the planned mix of similarity was not made compulsory, it was left to the discretion of the managers and supervisors to decide who they would hire—this influencing the extent to which they would follow what had been planned in the workshop. The actual number and amount of experience according to the level of similarity of the workers hired *after* the workshop is currently being assessed.

specialists to give temporary technical support as well as to provide novices with longer periods of on-the-job training.[27]

Three other sources of tacit knowledge have been used in the project in addition to hiring experienced people: *consultants*, *retired workers* and *vendors*. Thus, seven consultants with expertise in plant start-up and problem solutions in high-similarity industrial processes were hired to spend 262 person-weeks in the plant giving support to the Operations Team.[28] Subsequently, 22 retired workers from the most similar industry in Brazil are also currently being hired as process operational specialists to spend from 1 to 3 years in the plant from the hot commissioning phase onwards. Finally, technical assistance contracts are also being set up with vendors to establish their specialists in the plant for longer periods than is the standard procedure in industry. Bringing in these high-similarity specialists compensates for the low experience in some of the teams by increasing the mix of similarity in specific areas for a period long enough for lower similarity workers to go through the necessary (guided) experience.

'Lowering the water level': tension and discomfort

To use 'experience' with no qualifications is like having a lake with the water at its highest level, while the use of the levels of similarity reduces the water, allowing everyone to see the submersed mountains, peaks and depressions of knowledge and experience for themselves. 'Levels of similarity' has therefore created some discomfort and tensions in particular circumstances. Paradoxically, the reason why levels of similarity may create discomfort or tension is actually its main feature: It enables individuals, managers and organisations to become *aware of what they do not know.*

Thus, some people from the Brazilian company, who considered themselves experienced and were also considered by their colleagues to be so, soon realised that the many years of technical experience in the *mining* industry could not be directly translated into experience in the *nickel* industry. This caused some of the low-similarity white- and blue-collar employees to feel uncomfortable with the situation. However, while a few of them felt threatened and 'paralysed', others saw the situation as an obstacle that could be overcome with time within the current work and/or utilised it as an instrument to improve their team mix or to demand additional training for themselves. There were also instances witnessed during fieldwork where the realisation of their own technical limits actually caused low-similarity workers to be more careful with their evaluations and managerial decisions—and to seek technical support from their high-similarity peers when necessary.

[27] This was the strategy of the Operations Team from the beginning. Achieving the same for the Maintenance and Health & Safety Teams came about as a result of the discussion on the mix of similarity and the use of the Operations Training Programme as an example to be followed.

[28] Together, the three process engineers, one technologist and four process operational specialists from this consulting company have 237 years of high-similarity experience (an average of 26.3 years/person). The four process/operational specialists started their professional lives in low ranking positions—each with different jobs (welder, a product dryer operator, kiln operator and control room operator)—and finished as Senior Supervisor, Senior Process Specialist, Operation Advisor and Process Plant Superintendent, respectively. Together, these four people have 158 years of high-similarity experience (an average of 39.5 years/person).

Some tension also occurred when a few managers and supervisors who were hiring new employees had their decision challenged by colleagues or subordinates on the grounds that the candidate chosen was not of high similarity, as had been planned. Finally, there were incidents where they realised that they had not made a good selection for their team and there were only a few posts left to fill. In these cases, managers and supervisors were encouraged to provide their team with more technical visits and on-the-job training in other companies or to adopt the compensation strategy.

Discussion

Tacit knowledge situatedness

Managing tacit knowledge means having the activity or web of activities, those with experience in it and the embeddeding form of life as the starting points of any analysis. It is then necessary to identify the subject/activity under analysis before moving on to define who is of high, medium or low similarity. To omit circumscribing what is under discussion as a means of working with levels of similarity is the same as using 'experience' without qualification. In short, as with tacit knowledge, one person does not have high similarity *in everything*, but only *in something* connected to his or her previous life/practical experience.[29]

The situatedness or specificity for the application of collective tacit knowledge was noticed when establishing the standards of similarity during the workshop. The contrast between the distinct working experiences could only be achieved through refined *assessments* of *similarity*/*difference* and *relevance/irrelevance* made by high-similarity workers. This point was also clear when observing the differences in the participation of the high-similarity employees (HSEs) and the LSEs in two other meetings that had *risk assessment* as a central part of the activities.[30]

The strategic risk workshop During the first meeting, participants were asked to devise worst-case scenarios within the project and to assess their risk vis-à-vis the distinct types of losses they could lead to (i.e. financial problems as well as those of health, personal safety/environmental matters, reputation, social issues and regulatory/compliance difficulties). They were then asked to analyse the existing preventive

[29] Thus, a common mistake some low-similarity employees committed in the beginning was to pinpoint bad judgements or decisions made by those considered of 'high similarity' as the means to question the concept of levels of similarity. They did not realise that, although the ones pointed at had extended experience in the nickel industry, a close analysis of their working experience would show that they had no experience in the specific area/subject in which they had made the bad judgements or taken the wrong decisions. There was a case, for instance, in which a high-similarity specialist was criticised by decisions he took with regard to the slag dumping area, but this area did not even exist in the plants he had worked before!

[30] The meetings reported below occurred respectively in August and November 2010, and with the exception of the description of the latter, only minor changes were made to the submitted manuscript after September 2010.

and mitigation controls, to assess the residual risks and, if necessary, to propose new controls.

The HSEs came up with many worst-case scenarios based on their first-hand experience in previous plants or on stories from trustworthy colleagues; they were clearly able to defend the scenarios and quickly assess the probability of their risks within the project as well as to propose the necessary controls. During most of the technical discussions, the LSEs remained silent while a few HSEs would talk among themselves and would usually agree with each other's arguments. However, when the LSEs did come up with worst-case *technical* scenarios, these were not defensible and were mostly discarded by the HSEs. Nevertheless, when dealing with other cases—such as the likelihood of an indigenous protest closing the front gate of the plant and the possible losses—some LSEs were able to assess the risks, impacts and controls. In short, depending on the issue, the LSEs were no longer LSEs: Their previous experience was sufficient for them to fully participate in the workshop.

The manual slag tapping safety Hazop (hazard and operability studies) The goal of this meeting was to assess how risky the recently identified need to carry out manual slag tapping would be. As expected, a similar preponderance of the participating HSEs' arguments, cases, examples and proposals was noticeable. However, the interesting feature was that the two members of the Safety Team who were supposed to play a major role within the Hazop could not do so. Both were of low similarity and a lack of experience in dealing with safety *within the nickel industry* did not enable them to assess the task risks properly. Their main participation was then restricted to comments about the correct use of the Hazop *technique* and how the final result—a number within a risk matrix—was *meant to define* what could or could not be done according to the safety rules.

A strict interpretation and application of rules in the absence of experience was verified in other cases involving LSEs of the Safety Team. For example, there are compulsory standards for working within 'confined spaces', but the interpretation of a confined space for the safety team would not be the same as it would for high-similarity workers. Hence, a low-similarity safety worker would stop bricklayers from working inside the electric furnace because they had not had the training concerning a 'confined space', even though the walls of the furnace were still being constructed with no risk of there being poisonous gases inside! Other cases were cited by HSEs, including one in which the strict application of the rules actually led to an unsafe situation.

The pattern of interpreting and applying rules and methods strictly was not a privilege of the Safety Team, however, but a common situation when those judging and making decisions were of low similarity. In addition, a causal relationship was noticed between the workers' experience in the field (i.e. levels o similarity) and the adopted—prescriptive versus practice-based—approach.

Making correct judgements and the two approaches to learning and doing

Low-similarity workers who cannot judge similarity, relevance or risk properly have a tendency to adopt a prescriptive approach to *doing*. They apply rules to the letter

even though this may be unproductive, uneconomic or even dangerous. They believe in management methods and techniques as means to achieve good results. Many discussions took place during fieldwork in which LSEs cited rules or praised the use of methods and standard matrixes in order to back up their position while high-similarity workers recalled and resorted to experiences in previous nickel plants to back up theirs.

Low-similarity workers are afraid of being accountable for problems in the future, and for them, the statement (sometimes explicitly spoken) 'I am/was following the standard/method' is a safe alibi. It is as if they were 'working-to-the-rule' without being aware of it. Of course, when they were novices working for high-similarity supervisors or engineers, the latter would simply start teaching them how to apply the rules sensibly, even breaking or creating new ones when necessary for production or quality purposes, but without compromising safety. Problems came about, however, when the LSEs had more power than their high-similarity colleagues (or subordinates) and the final decision was theirs.

Differences in judgements have therefore led to some clashes between the high and low-similarity employees as regards technical issues or managerial decisions with a bearing on technical points. Thus, 'off the cuff' comments such as 'they have no idea about operating a plant like this' or 'they will learn when the operations start' were heard in circumstances where HSEs felt that their views about what they considered to be serious problems were not taken into account by their low-similarity colleagues or superiors. In general, the complaints were that the strict application of rules and methods (some of them typical of the mining industry) and decisions not based on nickel production experience were reducing the overall productivity of the project, producing delays and, in some cases, even creating situations that should have been avoided.[31]

A prescriptive approach to *learning* was also spotted with regard to the application of a set of corporate training programmes on safety. The rule was that each employee should go through one or more form of safety training depending on the activities each had to perform. For instance, if the person had to work within confined spaces or high up, the respective courses should be taken. The problem is that those who did the first evaluation were of low similarity. They did not know exactly which types of activities were performed by each function and how this could change according to the adopted work organisation. The outcome was therefore a conservative one: to train as many people as possible in almost all the existing courses. This led, however, to such a high amount of training that it generated a complaint about the employees being taken from relevant tasks. A new assessment was then made by HSEs leading to a considerable reduction in the number of hours of training. For instance, in the Operations Team, there was a reduction of 66% in the training workload—from

[31] Moreover, remarks were made about the reluctance of low-similarity workers leaving their desks and going to the construction site on a daily basis or in order to analyse the circumstances or problems being discussed *in loco*. In the words of an HSE, one of the final outcomes of the company change of fields is that 'after the ramp-up our [production] cost will be two times the cost of [the name of the other older nickel company in Brazil]'. The position held by HSEs with regard to the future can then be roughly divided in two: There are those who believe that the project will fail to be as profitable as it could be or, given the lower margins in the nickel industry compared to the mining industry, that the project will *have* to adapt itself to the form of life of producing nickel.

40,497 h (1,657 courses) of training before the revision to 13,813 h (693 courses) after it—which meant an estimate saving of approximately US $171,500.00/year.[32]

The premise used by the person who did the new evaluation for the Operations Team is the relevant point in this case—in addition to the judgemental ability of the HSEs already discussed. He explained to me that, given the internal rules, those who attended the course are considered by the company as being formally qualified to do the job. Nevertheless, based on his experience (or on what we called the practice-based approach to learning), classroom training is insufficient; a person would only be prepared to work safely after working under the guidance of experienced peers for a certain amount of time. In other words, an inexperienced but formally qualified person was viewed by him as being something dangerous. Thus, to train a group of people who would not have the opportunity to work with the subject afterwards would actually increase the risk of accidents. In sum, reducing the training meant reducing the risk in this case! Nevertheless, this is only because the rule for 'qualifying' someone to do a job (in this and in most of the companies) is based on a prescriptive approach to learning: to attend a classroom course. If the criteria were based on the practice-based approach, only more training would increase safety!

Judging within stable versus frontier technologies

A parallel can be made between 'producing nickel' and the ability to judge within stable and frontier technologies. In the case of the Challenger explosion, for instance, Collins and Pinch (1998) argue that there was no way of avoiding the accident because NASA and its affiliates were dealing with a frontier technology, where the parameters of what is right or wrong were still being agreed upon. That is, what was considered to be an 'acceptable risk'—to use the O-rings in low temperatures—proved not to be so after the launching. To think there was a way of avoiding the explosion is a 'retrospective illusion'—that is, to analyse an event or decision in the past with the knowledge you have nowadays (Collins and Pinch 1998). This means that stable technologies can be defined as those technical devices for which there is an overall agreement with regard to the correct judgement of most of its components and their way of functioning independently and within the system.[33]

Nickel production is not similar to the Challenger case. It is a stable technology where the ways of going on are already established and agreed upon, but this is true only for *those who already belong to this technical culture*. The problem is that a technology that is considered stable may become a frontier technology in certain situations. One of these is when companies enter into new domains and their employees are asked to *learn* how to operate and maintain the associated technology. Another is when HSEs are asked to deal with new technologies they have never worked with. Although the employees in both cases may be very experienced, hard-working and committed, they are immediately turned into low-similarity workers in the

[32] Exchange rate of 10 July 2010. The calculations were made by a specialist engineer, and the final numbers were taken from the report to his manager.

[33] Thus, when there is a failure in a technical device that nobody can understand, stable technology becomes frontier technology until a final verdict on what caused it is agreed. This is illustrated by the many debates on the reasons for the falling of the Air Bus 330-200 airplane that was travelling from Rio to Paris in 2009.

new domain or concerning the new technologies, therefore becoming subject to the risks commonly associated with actual frontier technologies. In sum, if the perspectives of those *who are learning* are adopted, any mistakes made can only be defined as such a posteriori, when those learning actually become aware that their actions were a 'mistake'. This has consequences for the way mistakes and errors are treated within organisations.

It is widely recognised that there is a learning curve when companies decide to work with frontier technologies or enter into a new field. The problem is that companies frequently do not realise or properly assess the differences between what they are used to doing and what they are about to do. In these cases, they are not prepared for or even aware of the *learning cost* that can be derived from the *learning curve*. Hence, the tendency is to blame individuals when the problems are easily 'identified' *afterwards* but are not linked to the learning phase and, by extension, to the decision made.[34]

The only way to minimise the risks of having 'stable' technologies turned into 'frontier' ones is to make a proper mix of low and high-similarity workers and to leave the technical decisions to the latter group. The paradox is that high similarity is necessary in the (new!) area in order to be able to judge what a 'proper' mix is and to actually obtain it. In order to avoid such a paradox, a similarity-based selection system should be applied to the hiring of all the employees, *from the shop floor to top management* and *from the very beginning* of every new project.

Explicit knowledge?

Two definitions of explicit knowledge were presented at the beginning of the paper. In one, the idea of explicit knowledge is not seen as problematic since 'this kind of knowledge can be transmitted across individuals formally and easily' (Nonaka and Takeuchi 1995, p. viii) and another in which 'all knowledge is tacit or rooted in tacit knowledge [and the idea of a] wholly explicit knowledge is unthinkable' (Polanyi 1969, p. 144). A third thought-provoking and more heterodox view will be put forward here.

There is no such a thing as 'explicit knowledge'. Books, standards, machines, signs, mathematical models, expert systems, smoke signals and so forth are neither 'explicit' nor a 'type' of knowledge. A thought example clarifies the point. Imagine that you ask someone in the street where the public library is and you are able to find it based on the person's answer, 'take the first right, it's the second building on your left'. In physical terms, the person has produced 'sound waves'. It is misleading to think that they contain something *within* them; they are just sound waves! So, where is knowledge here? The knowledge is in the person who knew where the library was, since he or she was/is a member of a form of life of which reading, books and library buildings form a part. It is also present in the participants of the dialogue who have the practice of both producing and attributing meaning to specific sound waves the

[34] A Brazilian expression—'engineer of a finished construction'—was heard a few times from those in charge of particular tasks as means of defending themselves from the 'wise after the event' approach adopted by others when problems were detected only later on, such as at the end of assembly or beginning of operations.

use of which is shared within those called English speakers (Wittgenstein 1976 [1953]).[35] Thus, you learned where the library was *from* the person *through* the sound waves. Similarly, books are just marks on a paper—a physical medium *through* which (some) readers may learn *from* their authors. They have no knowledge 'inside' them. The problem is that books are so close to us that we overlook the fact that they are like sound waves. In short, it only makes sense to say something is 'explicit' if we include it with the answer to the question: explicit for whom, when and in which social group? Thus, as the conversation with the experienced supervisor has shown, an operational standard requires previous experience for its user to break the regress of rules with regard to its interpretation and application to real situations. The same pattern was noticed with regard to high-similarity employees applying rules sensibly while their low-similarity counterparts do it rigidly and clumsily.

Whatever has been called explicit knowledge is better understood as 'products of knowledge' (Ribeiro 2007a), that is, any*thing* that can be used by enculturated actors or can be incorporated in machines, software and so forth. Knowledge is to be found in those who write books and design technical artefacts or systems as well as in those who are able to use, maintain, update or improve them accordingly. Products of knowledge are just reifications of 'bits' of the much more lively, ever-changing and ungraspable human knowledge—thus, the need for designers and manufacturers to update a software or upgrade a technical artefact as practices changes with time. At best, it may be said that they offer opportunities 'for use' or 'for learning' from their writers and designers, but these are dated and only available to enculturated actors.[36]

Moreover, *real* learning and *proper* use only happens after action is taken based on what is thought to be understood from reading, doing reverse engineering or by trial-and-error and its outputs are properly assessed. To be 'explicit' is therefore not an attribute of the written word or of technical artefacts. It is rather the outcome of the *moment* in which full-blown members of a form of life deal with them successfully. If this argument does not help to discontinue the use of the concept of 'explicit knowledge', at least it will highlight how misleading it is and the relevance of focusing on tacit knowledge development within organizations.

Research caveats

This ongoing action research illustrates that it is possible to manage tacit knowledge. This is not the same as saying that the Brazilian company will succeed in its objective of having a safe and smooth start-up of the nickel plant. Because the research focused on two areas, it did not comprise an overall analysis of how many managers and supervisors within the other areas of the plant tried to implement the proposed methodology and how many of those who tried actually succeeded. More importantly, since it was being applied for the first time, the methodology itself has since been subjected to analysis and refinement. With regard to the levels of similarity, for instance, there is no guarantee that the mix of similarity planned in the workshop—even if achieved—will be the right mix.

[35] If the bilingual person had produced sound waves equivalent to the written instruction 'vire primeira à direita, é o segundo prédio à esquerda', you would probably still be trying to find the library!

[36] It goes without saying that full enculturation is not even a guarantee of 'explicitness', as the 'death of the author' discussion illustrates. I thank Francisco Lima for drawing my attention to machines as having 'opportunities for use'; this is where the analogy to books came from.

For a sociological analysis of how to make a proper mix, it would be necessary to make a study of the contributions and differences in performance between medium and high-similarity personnel working in different teams and functions *after the plant start-up*. It is also essential to understand what medium and high-similarity workers can and cannot transfer from their previous experiences to the nickel project and the reasons underlying both situations. Finally, it is important to gauge the time it would take for medium and high-similarity workers to adapt to the specifics of the nickel plant—and to decide when this point has been reached.[37] These research caveats are not a problem, however. Instead, they indicate that a new area of interdisciplinary enquiry has been opened and more research on these topics is necessary.

Conclusion

All types of tacit knowledge are necessary for producing nickel smoothly and safely. During the pre-operational phase, however, the presence of collective tacit knowledge emerges more prominently: The ability to make correct judgements is at the heart of deciding how every single part should work in the future plant. As the result of the interplay between the theory and fieldwork, three *types of judgement* were put forward and verified in a variety of meetings, some of which were here discussed. This has contributed to and endorsed the theoretical discussion on what collective tacit knowledge actually refers to and enables one to perform. However, as any type of tacit knowledge is developed only through immersion within a form of life, a closer look at what experience one is referring to is necessary.

The concept of *levels of similarity* has been proposed to qualify experience in theory and in practice. It facilitates the identification of *those who possess* the tacit knowledge specific to a given context by differentiating between the previous experiences of experienced professionals vis-à-vis the activities they will perform. This is possible when functions that were once considered as a 'whole' have their specifics separated by those who can judge similarity/difference and relevance/irrelevance properly in each case. This highlights the 'situatedness' of tacit knowledge or, in other words, the connection between the practices within which distinct types of tacit knowledge are developed and the ones to which they can be (successfully) applied later on.

To qualify existing or desired experience leads to a myriad of ways to manage tacit knowledge, two of which deserved our attention. One is the possibility to identify the relevant tacit knowledge and plan its distribution within the existing working groups through better mixes of similarity. This guarantees the presence of high-similarity employees in all areas of the plant—thus contributing to company goals; at the same time, it also provides the conditions for tacit knowledge transfer between these

[37] In addition, it would also be misleading to say that this intervention in the nickel plant will be *the* reason for a possible successful start-up or, on the contrary, that possible future failures, delays or accidents will be *solely* due to the fact that the right mix of similarity was not achieved. Given the plant's complexity and history, there are other parties (e.g. plant designers, vendors, builders) and factors (e.g. the market, changes made in the project, budget cuts, interests, climate) impacting on what will or will not happen during and after the start-up.

employees and the local novices. Second, Human Resources practices focused on tacit knowledge enhancement and recognition (e.g. similarity-based selection, hiring and attraction/retention systems) can now be thought of with the objective of achieving a tacit knowledge sustainable project in the shortest time possible—with positive impacts on its economic and social sustainability.

Two opposing *approaches to learning and doing—prescriptive* versus *practice-based*—were noticed and connected to the distinct levels of similarity. A prescriptive approach to learning was shown to create an illusion—for trainees and those who believe in it—that people are qualified to do the job, when actually they are not. Formal training in classrooms is only the beginning of an enculturation process that usually takes years to complete and only succeeds if based on a practice-based approach to learning—such as a structured on-the-job training. With regard to doing, it was seen that those with low similarity usually stick to a prescriptive approach which leads to uneconomic, inefficient and even dangerous situations. On the other hand, speed in acting with autonomy and assurance is the trademark of high-similarity employees. Due to their judgemental abilities, they are able to attribute meaning to rules or operational standards and apply them in the light of the specifics of the situation. This makes rules and standards 'explicit' for them, while they may be only marks in a paper for some novices.

In sum, a conceptual framework that enables managers and practitioners to think—and act—in terms of tacit knowledge transfer and development has been proposed and applied to a real case: a case of large-scale enculturation. Employees with no previous history in the nickel industry are subject to different kinds of experiences as part of their pre-operational training and daily work. Moreover, workers, consultants, retired individuals and vendors who possess previous experience in the production of nickel or with similar equipment/processes are also hired. Their major role is to (re) create a new form of life within the Brazilian firm, the one of producing nickel safely and efficiently. This has required an analysis of the types, amount and the distribution of tacit knowledge among the workforce and the establishment of the conditions for its development by non-enculturated workers.

This case indicates the feasibility of designing and improving policies, management systems and human resources practices focused on the creation, development and maintenance of the distinct bodies of tacit knowledge among the workforce. The result tends to be the capitalisation of the workers' experience and associated tacit knowledge—translated into better operational, safety, societal and, of course, financial outcomes—in addition to their formal recognition by and within organisations.

Acknowledgments This paper would not have been possible without the full support of the Brazilian company where the fieldwork took place as well as the support of its employees; I am enormously grateful to them for allowing me to share their professional experiences and donating much more of their time than would normally be expected. I am also indebted to Fernando Campos Guimarães, the person who first believed in the research and has been fully supportive throughout – in addition to Thiago Souza Magalhães who first introduced us. My very special thanks also go to João Gilberto Queiroz, who has been an outstanding research gate keeper and has also helped to make concrete the idea of tacit knowledge management inside the nickel project – it was his idea and initiative to hold the Operational Readiness Workshop. I am further indebted to Guilherme Matosinhos, a Production Engineering undergraduate student who acted as my research assistant for 18 months and to two other students, Virgínia Giani Casagrande and Bárbara Feitosa, who voluntarily helped with specific tasks. My thanks also go to Raoni Rajão and Ödül Bozkurt, to Harry Collins and Rob Evans, and to Evan Selinger, David Stone and Greg

Feist who have all invited me to present seminars about the ongoing research. These were held at the Department of Organisation for the Work Seminar Series 2009/2010 at Lancaster University, the III Workshop on Studies of Expertise and Experience (Cardiff University, November 2009) and at the National Science Foundation Workshop 'Acquiring and Using Interactional Expertise' held in Berkeley, USA, August 2010. The comments, criticisms and suggestions coming from these audiences were of great value as well as those from people who read previous drafts of this paper, namely, Raoni Rajão, Francisco Lima, Stuart Dreyfus, and Laura Cançado Ribeiro – to whom I am also grateful. Finally, a special thanks to Kay Shipton for her outstanding proofreading work.

References

Blackler, F. (1995). Knowledge, knowledge working and organizations: An overview and interpretation. *Organization Studies, 16*(6), 1021–1046.

Bloor, D. (2002 [1997]). *Wittgenstein, rules and institutions*. London: Routledge.

Collins, H. (1989). Computers and the sociology of scientific knowledge. *Social Studies of Science, 19*, 613–624.

Collins, H. (1992 [1985]). *Changing order*. Chicago: The University of Chicago Press.

Collins, H. (2001). Tacit knowledge, trust and the Q of sapphire. *Social Studies of Science, 31*(1), 71–85.

Collins, H. (2007). Bicycling on the moon: Collective tacit knowledge and somatic-limit tacit knowledge. *Organization Studies, 28*(2), 257–262.

Collins, H. (2010). *Tacit and explicit knowledge*. Chicago: The University of Chicago Press.

Collins, H., & Pinch, T. (1998). *The Golem at large*. Cambridge: Cambridge University Press.

Cook, S. D. N., & Brown, J. S. (1999). Bridging epistemologies: The generative dance between organizational knowledge and organizational knowing. *Organization Science, 10*(4), 381–400.

Dreyfus, H. (2009) *On the internet*. (2nd Ed.) London: Routledge.

Dreyfus, H. L. (1979 [1972]). *What computers can't do: The limits of artificial intelligence*. New York: Harper and Row.

Dreyfus, H. L., & Dreyfus, S. E. (1988). *Mind over machine—the power of human intuition and expertise in the era of the computer*. New York: Free.

Gherardi, S. (2000). Practice-based theorizing on learning and knowing in organizations. *Organization, 7* (2), 211–223.

Gorman, M. E. (2002). Types of knowledge and their roles in technology transfer. *The Journal of Technology Transfer, 27*, 219–231.

Kuhn, T. S. (1996 [1962]). *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. New York: Cambridge University Press.

Leith, P. (1986). Fundamental errors in legal logic programming. *The Computer Journal, 29*(6), 545–552.

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company*. Oxford: Oxford University Press.

Nonaka, I., & Von Krogh, G. (2009). Tacit knowledge and knowledge conversion: controversy and advancement in organizational knowledge creation theory. *Organization Science, 20*(3), 635–652.

Polanyi, M. (1969). The logic of tacit inference. In M. Grene (Ed.), *Knowing and Being: Essays by Michael Polanyi* (pp. 140– 44. Chicago: University of Chicago Press.

Polanyi, M. (1983 [1966]). *The tacit dimension*. London: Routledge & Kegan Paul.

Ribeiro, R. (2007a). *Knowledge transfer*. Unpublished doctoral dissertation, Cardiff University, School of Social Sciences, Cardiff, UK.

Ribeiro, R. (2007b) The language barrier as an aid to communication. *Social Studies of Science, 37/4*, 561–584.

Ribeiro, R., & Collins, H. (2007) The Bread-making machine: Tacit knowledge and two types of action. *Organization Studies, 28/9*, 1417–1433.

Tsoukas, H. (2005) *Complex knowledge*. Oxford: Oxford University Press.

Winch, P. (1990 [1958]). *The idea of a social science and its relation to philosophy*. London: Routledge.

Wittgenstein, L. (1976 [1953]). *Philosophical investigations*. Oxford: Blackwell.