Learning Groups: What Types are there? A Theoretical Analysis and an Empirical Study in a Consultancy Firm

Irma Bogenrieder and Bart Nooteboom

Organization Studies 2004; 25; 287
DOI: 10.1177/0170840604040045

Learning Groups: What Types Are There? A Theoretical Analysis and an Empirical Study in a Consultancy Firm

Irma Bogenrieder and Bart Nooteboom

Abstract

This article attempts to unpack the notion of 'communities of practice', in more detail than has been done before, and looks more generally at intra-organizational groups for learning. First, it gives a theoretical analysis of the relevant characteristics of learning groups which affect the possibilities and conditions for the sharing and joint development of knowledge. These characteristics include opportunities for learning (on the 'competence' side of relations), relational risk (on the 'governance' side of relations), and the effects on both from the 'structural features' of groups. On the competence side, it analyses the implications of different types of knowledge and learning, and the trade-off between stability of relations (for the sake of mutual understanding and trust) and flexibility of relations (for the sake of variety as a source of learning). On the structural side, it considers the effects on competence and governance of network density, the strength of ties, structural holes, and stability of group membership. On the governance side, it considers psychological risk, in loss of reputation or legitimation, career risk and risk of competition, and risk of lock-in into the group. Trust yields one basis for dealing with such risks, and the article discusses what that means and how trust develops. Next, the analysis is used in an exploratory empirical study of a consultancy company, to see if the theoretical framework can explain the occurrence, structure, functioning, and performance of learning groups found in practice.

Keywords: organizational learning, communities of practice, learning groups, teams, social networks, competence, governance, trust, consultancy firms

Irma Bogenrieder
Erasmus University Rotterdam, The Netherlands

Bart Nooteboom
Erasmus University Rotterdam, The Netherlands

Organization Studies
25(2): 287–313
ISSN 0170–8406

There are different levels of learning in organizations. On one level, people learn individually by adopting knowledge from others, that is, 'knowledge sharing'. On another level, people jointly develop knowledge that is new to an organization. The literature on knowledge management tends to focus only on knowledge sharing. That can be counterproductive, because measures taken to support knowledge sharing may impose uniformity of meanings and procedures that limit the variety needed for the development of new knowledge. Thus, we need to look at both knowledge sharing and knowledge production.

An issue in the theory of organizational learning is the relation between learning by people and learning on the level of the organization (Cook and Yanow 1996; Weick and Westley 1996). The link between the two levels

DOI: 10.1177/0170840604040045

arises in teams, groups, or 'communities', such as 'communities of practice'. Wenger and Snyder (2000: 139, 140) characterize a community of practice as a 'group of people informally bound together by shared expertise and passion for a joint enterprise', which can 'drive strategy, generate new lines of business, solve problems, promote the spread of best practices, develop professional skills, and help companies to recruit and retain talent'. On the one hand, this characterization is very broad, with a wide scope of activities. In our view, there is a need to specify different communities, with different aims and features. On the other hand, the characterization is narrow, in the sense that learning is connected with 'shared expertise' in a 'joint enterprise'. We want to cast our net wider than such communities, which is why we employ the wider term 'learning groups'. Thus, the question we address in this article is what different learning groups there may be (in a theoretical analysis) and to confront this with what we actually find (in an empirical exploration).

We propose that for the study of learning groups we need to consider three related aspects. First, we need to consider the 'technical' or 'competence' side, in terms of types of knowledge and learning. Second, we need to consider the motivational or 'governance' side: how to overcome obstacles that may prevent people from opening up to the sharing of knowledge. Third, we need to include structural features of groups, because, as we know from social network analysis, they affect both competence and governance.

The article proceeds as follows. In the first part of the article, we turn to theory. For the competence side, we discuss the theory of knowledge and learning and the corresponding features of learning groups. Next, we turn to structural features that affect competence and governance. For the governance side, we turn to issues of motivation and relational risk. This yields an inventory of relevant features for learning groups and a set of hypotheses concerning how these features affect each other and how they affect the performance of learning groups. In the second part of the article, we turn to the empirical exploration of learning groups in a consultancy firm. The aim is partly exploratory (to see what learning groups we find) and partly confirmatory (to test whether our theoretical framework can explain how they work, or fail to work).

Theory

Incommensurable Paradigms?

In our analysis, we employ a social constructivist theory of knowledge and learning. The basic idea is that people perceive, interpret, understand, and evaluate the world according to mental categories that they have developed in interaction with the physical world and with others. This view entails that knowledge cannot be claimed to be objective and that cognition is to a greater or lesser extent idiosyncratic, since it depends on a variety of experience from which it is built up. Yet, in our view, this does not necessarily imply that any

view is as good as any other. While people may have contrasting and at times irreconcilable views of the world, so that they cannot always agree on 'the facts', they can often agree on at least some shared meanings and observations (that is, agree on some jointly accepted 'facts') and enter into argument about whose ideas best fit them. One can agree on 'surface' elements of theory, even while strongly disagreeing about 'deep' fundamentals. In other words, while sometimes paradigms can be irreconcilable, often people can successfully set up a debate from different perspectives and combine elements from different perspectives in a coherent fashion. As a result, we claim that a constructivist view does not necessarily lead to radical relativism. In other words, some form of positivism can be combined with some form of constructivism.

While we fundamentally disagree with the perspective of transaction-cost economics (TCE), we maintain that the notion of 'specific investments' that people may need to make in relations, which make them mutually dependent and entail a relational risk of 'hold-up', remains useful. In fact, a constructivist view adds to the notion of such investments: because they see the world differently, people often need to build up mutual understanding, in a shared language, before they can share or jointly develop new knowledge. Such investment may be highly specific, particularly when the knowledge involved is highly tacit. Some critics would say that transaction-cost thinking and a constructivist, interactionist view constitute irreconcilable paradigms. That may be the case. However, from transaction-cost economics we adopt only the single notion of lock-in due to specific investments. This can be done without buying into the fundamental behavioural assumptions of that theory, which we reject. The notion of specific investments does not by itself entail any behavioural assumptions. We have fundamental objections to the foundations of TCE. We claim, however, that the notion and consequences of specific investments can without inconsistency be built into our constructivist, interactionist view of learning, in which, counter to TCE, trust and learning play a central role.

Theory of Knowledge

In their account of communities of practice, Brown and Duguid (1996: 60) employ an 'activity-theory' of knowledge (see, for example, Blackler 1995), in which action and learning feed each other, and they view 'learning as a bridge between working and innovation'. They employ the notion of 'canonical' and 'non-canonical' or 'procedural' (Cohen and Bacdayan 1996) knowledge. Canonical knowledge entails decontextualized, codified, and formalized rules for operation. Inevitably, such rules cannot cover the richness and the variability of practical contexts. It is by context-dependent deviations from canonical rules, with the ensuing need for improvisation and experimentation (Brown and Duguid (1996) employ Levi-Strauss's concept of bricolage), that learning arises in interaction between members of the community. This is based on 'storytelling' — to capture and share context-bound experience, to guide experimentation. As a result, communities emerge

from shared work practice rather than being designed *ex ante*. Edmonson (1999: 354) conceptualizes learning as follows: 'an ongoing process of reflection and action, characterized by asking questions, seeking feedback, experimenting, reflecting on results, and discussing errors or unexpected outcomes of actions'. Like Brown and Duguid (1960) and Edmonson (1999), and many others in this field, we also take learning to be a social process of interaction, and we employ an activity theory of knowledge. It is a truism to say that information is not the same as knowledge: to become knowledge, information needs to be interpreted in a cognitive framework. Activity theory descends from G. H. Mead's 'symbolic interactionism' in sociology and the view, taken from cognitive psychology, that intelligence is internalized action (Piaget 1970, 1974; Vygotsky 1962; Bruner 1979). Our view is also related to other 'constructivist', 'interpretative', or 'hermeneutic' views (compare Weick 1979, 1995). In contrast to the dominant 'computational representational' view in cognitive science, this leads to the viewing of knowledge in terms of 'situated action'. Knowledge and the meaning of words are not independent of context. They lie partly in the context of use and they shift from one context to another. This process precludes objective knowledge (or at least any certain knowledge whether or to what extent knowledge is objective).

Features of Knowledge and Learning

As noted before, learning groups can have a variety of knowledge and learning characteristics. Knowledge can have different *forms*: more or less tacit, canonical or procedural, context specific, and decontextualized. It can have different *contents*: professional expertise, skill, work perception and attitude, operation of projects, organization, markets (customers and competition), and 'meta-knowledge' on the location and reliability of sources of knowledge (Wegner et al. 1985). It can have a different *scope*: generic knowledge (beyond specific applications) or specific knowledge (for a given project or practice). The *aim* of learning may be joint production, problem solving, the development of new practices or products, exchanging experience from different projects, sharing codified knowledge, or the development of skills, training, attitude development, management development or organizational change. As noted before, it may be aimed at different *levels*: individual or collective learning. It may arise from sharing experience and knowledge in other ways than on the basis of shared practice in joint operations. In the constructivist perspective, there is more or less *cognitive distance* (Nooteboom 1999, 2000) between the people involved. In other words, it may be more or less difficult to understand each other. We will return to this later.

One definition of learning is the ability to respond differently to the same stimuli that obtained before, selecting from a given repertoire of responses. However, one might also learn to respond to new stimuli with new repertoires of action. This is related to the distinction between learning to do existing things better and to do new things. The first has variously been called first-

order, single-loop learning or learning for exploitation, and the second has been called second-order, double-loop or exploratory learning (Bateson 1973; Argyris and Schön 1978; Hedberg et al. 1976; Fiol and Lyles 1985). We call this the *depth* of learning. While we can make this conceptual distinction, in the process of learning the two kinds of learning do not stand apart from each other. Exploitation is based on exploration, and vice versa. We exploit what we have explored, and it is on the basis of exploitation that we explore (Nooteboom 2000). A central task of organizations is to find ways of combining the two. Nevertheless, groups can be primarily aimed at exploitation, which may then yield exploration, or at exploration, which should in due course yield exploitation.

Summing up, in knowledge we recognize the following features: content, form (degree of tacitness), and scope (generality versus specificity). In learning we recognize: aims, cognitive distance, level (individual or collective), and depth (exploitation or more exploration oriented).

Absorptive Capacity

Since mental categories have developed on the basis of interaction with others, in a sequence of contexts that make up experience, there will be 'cognitive distance' between people with different experiences, and cognitive similarity to the extent that people have interacted within a shared experience (though we do not wish to imply that 'cognitive distance' allows for any simple, one-dimensional scale). Cognitive distance yields both a problem and an opportunity. The opportunity is that we learn from others only when they see and know things differently. In the absence of claims of objective knowledge, interaction with others is the only path we have to correct our errors. The problem is that people may not understand each other and have to invest in understanding.

Tacit knowledge can never be fully expressed (let alone 'codified'): there is always some loss, due to disembedding from mental systems that are built on personal experience. In other words, knowledge is never identical between people. There are degrees to which disembedding takes place in communication. The 'storytelling' that was emphasized by Brown and Duguid (1996) carries more context specificity than communication of knowledge that is abstracted into canonical knowledge, where context specificity is shed. Expression by A can be absorbed by B only when B can fit it, or more precisely, reconstruct it, in his or her mental system, and vice versa. In other words, to communicate, A and B must have a mutual absorptive capacity that is sufficient for the task at hand. The more shared experience people have, the greater cognitive similarity will be, and communication can take place efficiently, in 'short-hand' communication with jargon that can be taken for granted, while not making sense to outsiders. Greater cognitive distance requires more effort to try and absorb what others do and say, and to communicate what one says and does in ways that help others to absorb it. That is why newcomers to the group have to start in 'peripheral participation' (Lave and Wenger 1991). This is needed to

develop shared tacit absorptive capacity by 'indwelling in the experiences, perspective, and concepts of other participants' (Von Krogh 1998: 114). As with the notion of 'tacit knowledge', the notion of 'indwelling' derives from Michael Polanyi (1962).

Structure

Structure is both the basis and the result of processes of interaction. The underlying intuition derives from the notion of structuration (Giddens 1984; Archer 1995), according to which structure enables (and constrains) action and action (re)constructs structure. In view of our 'situated action', interactionist view of meaning and knowledge, knowledge and learning are seen as 'embedded'. This is especially salient in the analysis of learning groups. Edmonson (1999: 351) proposed that 'to understand learning behaviour [in teams], team structure and shared beliefs must be investigated jointly'. We agree, and therefore we incorporate structure in our analysis. Here, we tap into the extensive sociological literature on network structure. There are three aspects: structure of the network, positions that people take in that structure (see, for example, Coleman 1990), and the strength versus weakness of ties (Granovetter 1982). One example of a position that someone takes relative to others lies in the notion of 'legitimate peripheral participation' (Lave and Wenger 1991). Another feature of position in a network is the degree of *centrality*: the degree to which an agent is connected to others who are not mutually connected. This has implications for power, in terms of access to alternative members, bargaining power, control of information and gossip, coalition formation, and a policy of 'divide and rule'.

Structure is *dense* when there are many direct linkages between all participants, and it is *sparse* to the extent that there are few such linkages. Related to this, *structural holes* (Burt 1992) refer to gaps in the network structure, with some participants isolated from other participants. There can be such holes within a group, and there are also external holes between groups. Burt argued that high density, with many direct ties, yields *redundancy*: there are many ties to maintain, while they yield little added value in access to new knowledge. Structural holes arise, in particular, between different groups. As a third feature, groups may be inclined to establish *closure* from outside influences (here, we understand closure to be the warding off of outside connections, rather than Coleman's (1990) notion of closure as full density, with everybody linked to everybody else). Important also, as a fourth feature, is the *stability* of group membership. Low stability entails frequent exit and the entry of 'outsiders'. A sixth feature is *structural equivalence*, that is, two or more group members who have ties to more or less the same members of the group. Structurally equivalent members, having the same pattern of ties in the group, may be rivals within the group.

In our view, the 'strength of ties' has four aspects. One aspect is *intensity*, which refers to the effort and commitment of resources involved, and to the scope of activities taken up in the tie (share of total activities). The resources that are committed are not necessarily only resources of money, time, or

effort, and may also include psychological resources (commitment, loyalty, fairness, and empathy). A second aspect is *frequency* of interaction, a third is *openness* of communication, and a fourth is *duration* of ties. Strong ties yield shared experience, which reduces cognitive distance. Durable ties enable the development of empathy and identification (McAllister 1995; Lewicki and Bunker 1996; Hansen 1999) as a basis for trust.

Summing up, we recognize six features of structure and position in structure: density, closure, centrality, stability, structural holes, and structural equivalence. We recognize four aspects of the strength of ties: intensity, frequency of interaction, openness of communication, and duration.

Relational Risk and Governance

Motivation concerns what one can gain from learning by interaction, but also what one can lose, in view of 'relational risk'. This yields the need for 'governance', defined as the mitigation of relational risk. In our theory of knowledge, we understand the notions of knowledge and cognition in a wide sense, encompassing perception, interpretation and evaluation, which includes emotion-laden value judgements. In other words, we see cognition and emotion (such as fear and suspicion) as linked (Merleau-Ponty 1964; Simon 1983; Nussbaum 2001). That is why it is essential to include emotional motivational factors.

Under the right conditions, people are willing and able to learn by exchanging, sharing, and jointly producing knowledge. The question now is what 'the right conditions' are. Edmonson (1999: 353) makes a distinction between the outcomes and the process of learning, and focuses on the process. We will look at both. Concerning the right conditions, Edmonson emphasizes the need for 'psychological safety', defined as 'a shared belief that the team is safe for interpersonal risk taking'. This is exemplified when in the exchange of knowledge, asking questions and so on, people place themselves at risk:

> 'for example, by admitting an error or asking for help, an individual may appear incompetent and thus suffer a blow to his or her image. In addition, such individuals may incur more tangible costs, if their actions create unfavourable impressions on people who influence decisions about promotions, raises, or project assignments.' (Edmonson 1999: 351)

Edmonson found empirical confirmation of the importance of psychological safety, and clearly this is an important aspect that needs to be included in any analysis of learning groups. What other 'right conditions' are there? Edmonson (1999) included not only the psychological risk of loss of face and reputation, but also 'tangible costs', such as risks to promotion, a raise, or referral for new projects. We propose an addition to the notion of psychological risk, and we explore in more detail the more tangible aspects of risk.

Our extension concerning psychological risk is as follows. There is not only a risk of loss of social legitimacy (loss of reputation and of acceptance by others) on the side of the one who opens up his knowledge, or lack of it, to others. There is also a problem in supplying knowledge that, rather than

exposing one's own weakness, on the contrary, yields highly advanced knowledge that is perceived as threatening by others, since it exposes *their* weaknesses. This entails the risk of becoming an isolated 'maverick' or outcast. Under such a threat, one may have to hold back one's knowledge, or carefully downplay it, and communicate it diplomatically in small doses.

The more tangible risks go beyond the results that loss of reputation may have for salary, career, or referral for new projects. We propose that it also includes two other features. One is risk due to 'spillover' of knowledge, defined as the sharing of knowledge that can affect competitive position. ('Spillover' is an unfortunate term, since it is too reminiscent of the old view, which we clearly reject, of knowledge seen in terms of a pipeline metaphor. According to that view, knowledge is a commodity that has an objective existence apart from people, and can be transferred like physical goods along 'communication channels'. However, it is a familiar term for competitive threat in knowledge sharing, and for the sake of recognition, we retain it.) The second feature is a risk of 'lock-in' in the group. We briefly discuss both.

Spillover risk entails the following. If the knowledge to be shared is close to one's 'core competence' or competitive advantage, then its adoption by others can threaten one's position to the extent that individuals compete for job positions and careers within the firm. Related to this, there is a risk of free riding. Colleagues may hold back on their knowledge while exploiting knowledge from others. Under a number of conditions such risk is absent or irrelevant (Nooteboom 1999). One such condition is that one's knowledge is so public that it will reach competitors anyway. A second is that the knowledge given is not close to one's 'core competence'. A third is that knowledge is so tacit that it does not spill over easily. It can be shared only in intensive collaboration. Related to this, a fourth condition is that potential competitors lack the necessary absorptive capacity. A fifth is that by the time potential competitors have absorbed and effectively implemented the knowledge for the purpose of competition, it has already changed. Lastly, if there is spillover risk, the value of knowledge one gains may exceed the disutility of spillover, so that one accepts spillover. If that is not the case, one way to control both psychological risk and spillover risk is to establish exclusiveness, by closing off communication with threatening others. This can be supplemented by the building of mutual trust. We will discuss the relevant aspects of trust later.

How about relational risk within groups? In our interpretation, communities of practice are characterized by the high density and strength of ties, with a certain durability and stability of relations. Spillover within the network is strong, and is, in fact, an essential part of its aim of sharing and jointly developing knowledge. This need not pose a threat, because a dense network yields density of communication (including gossip), as a basis for social control, in monitoring behaviour, a reputation mechanism, and the possibility of coalitions of some members to restrain others (Krackhardt 1999). Density and strong ties, with mutual dependence and specific investments, allow for close monitoring of conduct and the building of trust in mutual give and take ('reciprocity'). On the other hand, strong structural embeddedness

(Granovetter 1992) in a dense network with strong ties leads to cliques (Janis 1981). In time, that will yield limited cognitive distance, with the danger of 'group think' and lack of innovation, as Burt (1992) argued. Innovative but deviant ideas may be squelched. This will be the case especially when group membership is *stable*, with little entry and exit of members.

Structural holes arise, in particular, between different groups. They can be bridged by 'boundary spanners'. Burt (1992) suggested that such boundary spanners access new knowledge, providing cognitive distance, as a source of learning. However, a boundary spanner may be perceived as posing a threat in passing sensitive information from inside to outsiders. This can 'publicize' weaknesses, yielding loss of reputation. In other words, it increases the potential threat to psychological safety and career prospects. This may yield an argument for closure of the group to the outside world, as indicated before. Here, 'peripheral participation' may not be seen as 'legitimate', and new entrants are viewed with suspicion. However, the downside of such closure is that the scope for explorative learning and innovation shrinks, because there is a lack of renewal from new entrants in the group.

The risk of 'lock-in' results from investments in relations that are, to use that notion from TCE, 'specific' to the group or to people within the group. Here, the notion of specific investments takes on new forms, not recognized in TCE. As we indicated before, the sharing of knowledge requires appropriate mutual absorptive capacity. When this is not present, it has to be built up. Often this requires more or less intensive 'working together'. This applies more to the extent that knowledge is more tacit and context specific. When there is no prior trust in relations, that also has to be built up, in the form of personalized trust and on the basis of ongoing interaction. Both entail investments in group relations that are more or less specific. This can yield lock-in due to 'exit barriers': breaking relations and switching to others entails the need to engage in such investments anew. This lock-in entails that group members may be tempted to exploit each other, knowing that each will not easily exit from the relation. Also, people will want reassurance that relations will last sufficiently long to make specific investments worthwhile. Lock-in may not be perceived as a threat, but on the contrary, as a basis for psychological safety. On the other hand, one may arrive at the need to exit from the group in order to access alternative sources of learning not allowed by the group. One may want to exit, but lock-in then yields an exit barrier.

Summing up, we identify three types of relational risk: psychological risk, career and spillover risk, and the risk of lock-in due to relation- or group-specific investments in mutual absorptive capacity and the building of trust.

How can one 'govern' relational risk? In groups within organizations, formal control by legal contracts is not available, because group members do not have the requisite legal authority. Formal agreements, short of legal contracts, are not feasible when uncertainty is so large that they cannot be specified, as is typically the case in learning, especially explorative learning. Also, especially in learning, there may be no adequate basis for monitoring conformance to agreements. When, nevertheless, one pursues formal, detailed agreements, they can yield undesirable constraints on the scope of learning.

Part of the reason for this lies in the canonification of rules required in formal agreements, which cannot do justice to the richness and variability of emerging practice, as discussed by Brown and Duguid (1996). Furthermore, formal agreements can set in motion a vicious cycle of distrust, which may block the building of trust. Another form of control is hierarchical monitoring and intervention 'by fiat'. One problem with this is that, especially in learning, the 'boss' cannot adequately assess performance, if such assessment requires active participation in collegial collaboration. A third form of control is by means of financial incentives or penalties. However, this again requires reliable monitoring, which is often absent. A fourth way to mitigate relational risk is to establish mutual dependence, so that people need to reach a mutual reduction of relational risk. Here, people may be prevented from using their power over others by the fact that they are themselves dependent on them. Also, when people need each other without alternatives being available, are locked into each other, there is a greater pressure to develop trust. A fifth way is the use of a reputation mechanism. Here, people may behave well to protect their reputation, needed for future collaboration. As discussed before, relational risk can also be mitigated by group structure.

Trust

All the ways to mitigate relational risk discussed so far appeal to pure self-interest, and constitute forms of 'control' or 'deterrence' (Maguire et al. 2001). They yield extrinsic motivation that may 'crowd out' intrinsic motivation (Frey 2002). Often, trust does not only have extrinsic value, to mitigate relational risk. A trust-based relation may be valued as an end in itself. Trust is a complex subject which cannot be fully discussed here (for an systematic survey, see Nooteboom 2002). Here, we summarize some key features.

Governance by deterrence can be very costly, and can break down the basis for the development of personalized trust. Many authors feel that 'control', in the sense of deterrence, is foreign to the notion of trust, and that 'genuine', 'thick', or 'real' trust entails the expectation that others will not behave opportunistically even if they have both the opportunity and incentives for doing so. 'Real' trust is based on other, more social and personal foundations of trustworthiness. Social foundations may be found in shared norms of reciprocity or moral duty and obligation. This yields a sixth way to mitigate relational risk. Trust may also be more personalized, on the basis of empathy, identification, or friendship, which yields a seventh way. When trust is not in place prior to a relation, it has to be built up. Intensive collaboration can set in motion a positive cycle of emerging trust. Zand (1972) proposed a cycle in which trust engenders openness, yielding information, which provides a basis for the application and acceptance of mutual influence, which yields the willingness to demand less and accept more control from the partner, which further engenders trust. This is related to Hirschman's (1970) notion of 'voice'. When conflict arises, the first response is not to quit ('exit'), but to seek amends. One reports one's dissatisfaction, asks for an explanation, is

open to claims that one is oneself at fault, welcomes criticism, and asks for and offers help to 'work things out' together by solving problems, repairing shortcomings, and eliminating misunderstandings.

There are stages in the development of trust (Lewicki and Bunker 1996; McAllister 1995). In such development, we may reach the stage at which we can empathize with people, that is, get to know and understand their 'life world'. Here, and elsewhere, it is important to make a distinction between trust in competence and trust in intentions for fair dealing. Will a member of a learning group have useful knowledge to offer and be able to communicate it, and will he understand what we say (competence)? Will he not expropriate knowledge he receives and use it to compete, and will he, in fact, be committed to utilizing his competence for mutual benefit (intentions)?

Empathy helps in judging trustworthiness: it helps in attributing competencies, intentions, and honesty correctly. Beyond empathy, we can identify with people, to the extent that we share the same perceptions, interpretations and evaluations, in a common 'life world'. Here, we may not only understand, but also sympathize with weaknesses, and tolerate deviations from expectations. Empathy and identification are generally based on shared experience in the process of 'indwelling' mentioned before.

However, trust based on identification, friendship, or kinship can go too far, in unperceived relational risk. It may not be sufficiently robust under extremes of temptation or pressures of survival. Generally, unconditional trust is unwise: it is too much to expect that partners will be able to resist even the strongest temptations for opportunism or pressures of survival. For example, when conditions in a firm yield a threat of lay-offs, this may intensify rivalry in keeping one's job, which intensifies perceived risks. Therefore, it is generally wise to see trust as subject to tolerance levels: one will trust until events are perceived that exceed a partner's ability to resist temptation or pressures of survival. In routinized trust, when a relation has been going well for a time, one may no longer be attentive to opportunities or pressures for opportunism regarding oneself and others in the group. One reason why emotions are needed in cognition, as proposed before, is that they serve to trigger attention to tacit routines, catapulting them from 'subsidiary' to 'focal' awareness (Polanyi 1962) when threat arises.

Variables and Hypotheses

From the preceding analysis we now compose a list of variables that represent the relevant features of learning groups. They are specified in Table 1. Note that we specify 'potential' relational risks, and then look at governance to mitigate such risks. 'Financial incentives' are direct incentives based on some measurement of contributions to the group. They are not indirect incentives that might follow from project performance or increased knowledge or skill.

From our analysis we also derive a number of hypotheses, specified in Table 2. The hypotheses show a dilemma. On the one hand, for building mutual understanding and trust, in order to mitigate relational risk, one may require a dense network with strong ties. The investment required is largely

Table 1. Variables

Knowledge and learning	**Potential relational risk**
Content: professional, operation of projects, organization, market and competition, sources of knowledge	Psychological risk: loss of reputation, 'face', prestige, self-image
Distance: limited/large cognitive distance, high/low common absorptive capacity	Career and spillover risk: consequences of openness for rewards, career, perspectives for projects, and (internal or external) competition
Form: degree of tacitness	Lock-in risk: depending on relation- or group-specific investments in mutual understanding or the building of personalized trust
Level: collective/individual	
Scope: generic/specific	
Depth: exploitation/exploration	
Structure and strength of ties	**Governance (to mitigate relational risk)**
Density	Hierarchy
Closure, exclusiveness	Financial incentives
Centrality	Mutual dependence
Stability of group membership	Reputation
Structural holes	Dense network
Structural equivalence	Network position: centrality, spanning structural holes
Intensity of ties: degree and scope of mutual involvement	Shared norms of behaviour
Frequency of interaction	Routinization
Openness of communication	Competence trust
Duration	Intentional trust, on the basis of indwelling, empathy, identification

Table 2. Hypotheses

On cognition

H1 Learning by interaction requires intermediate cognitive distance: large enough to yield novelty and small enough to enable understanding. Understanding requires mutual absorptive capacity. Greater absorptive capacity permits greater cognitive distance.

H2 When cognitive distance is too large for mutual understanding, and absorptive capacities need to be developed, these have to be developed in interaction. This is the case especially when the knowledge involved is highly tacit. This may entail a specific investment, for which there must be a perspective for continued interaction to make the effort worthwhile.

On relational risk

H3 Knowledge exchange may entail three types of relational risk. One is psychological risk, in loss of reputation, loss of face, loss of self-image, or loss of legitimation by the loss of face one causes to others. The second is career risk, due to loss of reputation or legitimation and risk of spillover (free riding and competition). The third is lock-in risk due to specific investments in mutual understanding or the building of personalized trust.

On governance

H4 For learning by interaction, relational risk must be mitigated (eliminated or alleviated). Formal control of relational risk by formal agreements or by hierarchical control is not feasible to the extent there is uncertainty, and actions and results are insufficiently observable, as is often the case in learning. Other forms of control are financial incentives, reputation mechanisms, and mutual dependence. One may also mitigate one's risk by occupying a position of power in the group.

H5 Beyond such control, geared to pure self-interest, trust may serve to mitigate relational risk. It may be based on shared norms of behaviour or on personalized trust built up from collaboration. One needs to distinguish between competence trust and intentional trust.

On the effects of group structure

H6 Reputation mechanisms and the building of personalized trust are enhanced by dense network structure with strong ties.

H7 Ongoing strong ties in dense, closed networks lead to a reduction of cognitive distance, which enhances exploitation, but limits exploration.

Table 3
Community of Practice

Knowledge and learning
Content: operation of projects, product or process improvements
Distance: limited, with high mutual absorptive capacity, shared beliefs and norms, efficient communication in group-specific codes
Form: highly tacit, context-specific knowledge
Level: collective, for joint problem solving, shared tasks
Scope: specific knowledge, embedded in contexts of application
Depth: emphasis on exploitation, possibly as a basis for exploration

Structure and type of ties
Density: high
Closure: high, fairly exclusive relations, entry barriers
Centrality: low, with little hierarchy
Stability/volatility of group membership: fairly stable, but with some exit and entry of trusted specialists, via legitimate peripheral participation
Structural holes: low
Structural equivalence: low; complementary competencies
Intensity of ties: high
Frequency of interaction: high
Openness of communication: high
Duration of membership: more or less durable, ongoing

Potential relational risk
Psychological risk: high due to openness needed for intensive collaboration
Career and spillover risk: high due to intensive and frequent ties
Lock-in risk: high due to relation- or group-specific investments in mutual understanding and the building of personalized trust

Governance
Hierarchy: none
Financial incentives: none
Mutual dependence: high
Reputation: high
Dense network: high
Network position: limited centrality
Norms of behaviour: high
Routinization: high
Competence trust: high, due to intensive collaboration
Intentional trust: high, due to indwelling, empathy, identification

specific, which requires sufficient durability of relations. To limit relational risk with respect to other groups, one may need to close a group off from outside linkages. A dense network may also be needed for social control and a reputation mechanism. On the other hand, dense networks entail redundancy with respect to knowledge access. Durable ties lead to reduced cognitive distance, especially when the group is closed and has a stable membership. This may be good for exploitation, but yields a diminished scope for exploration (hypothesis 7).

Communities of Practice

To demonstrate our framework, we use it in an attempt to give a clear, detailed specification of the concept of 'communities of practice'. This is not a test of the validity of our framework, but of its analytical usefulness. For communities of practice, we propose the specification given in Table 3. The table reflects our interpretation of the concept found in the literature. If our interpretation is wrong, our framework may be used to indicate precisely where it is wrong.

Empirical Study

Methodology

Our empirical study of learning groups does not constitute a rigorous test of the hypotheses, with quantified measures and correlations of variables on the basis of a large sample. We propose that at this stage the analysis needs to be more exploratory and process oriented, using longitudinal case studies to study the structures and processes of learning in time. The aim is to see whether:

- The learning groups we find can be adequately characterized by the proposed variables (Table 1), and
- Their functioning and goal achievement can be explained by our hypotheses (Table 2).

Again, this is more a test of usefulness than of validity. Can our framework reconstruct and explain the cases that we observe?

Our case studies were performed in a consultancy firm that we name 'X'. There, we found four cases of groups that had explicit aims of learning. X is a global management consulting and information technology services company. The research was conducted at the Dutch arm of X, in the business consultancy unit, which consisted of about 80 people. In August 1999 and January 2000, 19 interviews were conducted at the Dutch headquarters of X with senior and junior business consultants and managing consultants. The managing consultants led functional groups of around 15 business consultants each. In some cases, we talked to the business unit manager. The interviews took about 1.5 hours and were recorded. We used 'triangulation' as a test of internal validity, by crosschecking accounts of each group made by different participants.

Case Studies

In firm X, there is strong emphasis on professional development and a drive for excellence. Knowledge workers are considered an asset. However, there is also a strong emphasis on individual responsibility for one's professional career and professional development. The general view is that the organization can facilitate professional development of the individual, but it is the knowledge worker himself or herself that must take initiative to do it. Individual consultants are also responsible for the acquisition of projects. As a result, in X there is considerable internal competition, yielding a potential spillover risk in knowledge sharing. Depending on the assignment, consultants work alone or in project teams with a client for between several months and several years. There is a large variation in the duration of project teams, and rotation of consultants across teams.

Consultancy skills fall into three domains: substance and content (that is, market development, tools, and methods), personal skills, and consultancy skills. On a more abstract level, the development of a consultant is described

within X as the development of 'being', 'skills', and 'knowledge'. Here, 'being' has to do with attitude and psychological resources concerning work, colleagues, and customers. Individual acquisition of codified knowledge is provided by numerous courses and training programmes. Diverse groups are experimenting with how to foster other forms of learning: development of tacit knowledge, skills, and 'being'. We identified four groups in which, according to the respondents, some type of learning takes place, or is intended to do so. For each group, we first give a brief description based on the taped interviews. We take the case descriptions recorded and indicate in italics the corresponding variables from our framework (Table 1), which are summarized for all cases in Table 4. To the description we add an analysis and make an attempt to explain what happens on the basis of our hypotheses.

Table 4
Variables of Groups

	Project team	Expert group	Professional development	Project improvement	Orientation
Knowledge and learning					
Content	projects	professional expertise	skills	projects	sources of knowledge
Distance	medium	limited	fairly high	high	high
Form	more tacit	less tacit	more tacit	more tacit	less tacit
Level	collective	individual	individual	individual	individual
Scope	specific	generic	generic	specific	generic
Depth	mainly exploitation	exploration	both	exploitation	either or both
Structure and type of ties					
Density	high	low	high	low	low
Closure or exclusiveness	medium	low	high	low	low
Centrality	low	medium to low	low	fairly high	low
Stability of membership	high	medium	medium	low	low
Structural holes	low	high	high	high	high
Structural equivalence	low	high	low	low	low
Intensity of ties	high	low	high	low	low
Frequency of interaction	weekly	high	low: monthly	low: monthly	?
Open communication	high	limited	high	low	low
Duration of ties	medium	various[a]	medium	low	low
Potential relational risk					
Psychological	high	limited	low	high	low
Career and spillover	high	high	low	high	low
Lock-in	high	low	medium	high	low
Governance					
Hierarchy	none	none	none	some	none
Financial incentives	none	none	none	none	none
Mutual dependence	high	high	high	low	low
Reputation	high	high	limited	low	low
Dense network	high	limited	limited	low	low
Network position	?	?	?	?	?
Norms of behaviour	high	high	high	low	low
Routinization	high	low	high	low	low
Competence trust	high	high	high	low	low
Intentional trust	high	low	high	low	low

Note:
[a] Some of this type of group had more, and others less durable ties.

Since we did not delve into the details of the network position of individual group members, we have no observations on network position. Note that hierarchical control is hardly ever used for governance and that direct financial incentives are never used. This is related to highly voluntary participation and self-initiative to engage in learning, with only indirect incentives deriving from improved performance or competence.

Project Teams

Project teams consist of several consultants — sometimes drawn from other business units within X (*cognitive distance*). The aim is good project performance. Thus, the level of learning is *collective*, and the scope is project *specific*. In order to facilitate and encourage cooperation within a project team, X had recently begun to experiment by paying special attention to team processes. We investigated an attempt that was made to find a balanced composition for the team with the help of Belbin's (1984) questionnaire (*network structure*). Before the start of the project, the members of the team were given a few days off in order to get to know each other (*dense network*, *trust building*, and *building mutual absorptive capacity*). This was experienced as very valuable for team members to explain their personal situation. This information helped to reduce behavioural ambiguity and to develop *shared norms of behaviour*. There was no change in membership during the project (*stability* and *closure*). At the time of the present research, this project finished. The members of the project team did not believe that the common practice evident during the project would be continued as a community of practice (*limited duration of ties*).

The team members reported that they learned a lot from each other, especially *tacit knowledge*. For example, the team used a certain tool that all members had learned to use during official courses organized by X. One team member reported that the tool, which was used by the team, is now much clearer than before this project: 'Of course, I already knew the tool, but now I have seen how my colleagues apply this tool' (*competence trust*). Asked why he now understands this tool much better, he said that he has now really experienced how this tool is applied.

The team also developed new methods for implementing changes in the client organization. Asked for the reasons why the project members were inclined to try new methods, the answer was that they developed these methods together and they have committed themselves (*mutual absorptive capacity*) in *intensive ties*. As one respondent said: 'The manner in which the workshop was set up was decided by the group as a whole, and things were discussed there; I would never have come up with it on my own' (*dense structure* and *no centrality*). We conclude that the team was primarily oriented toward *exploitation*, but collaboration in practice yielded some *exploration*.

The basis for interaction can be described as indwelling, where members really delved into a colleague's use of a method. There was a lot of mutual adaptation and voice:

'Why I started applying things I hadn't before — because we had agreed to do it this way; you know, you don't have to agree with everything, but if you are going to work

as a team [*intentional trust*] ... We had decided that we would do it in this fashion and we did it that way, and I did a number of things I normally never would have done [*empathy*].'

There was a team meeting every week (*frequent interaction*), in which both the progress made on the project and the team process were discussed. During these moments of reflection, possible conflicts or dissatisfaction within the team were discussed in a straightforward and timely fashion (*openness of communication*). The attention given to the team process was experienced as very valuable. It supported the development of a joint enterprise:

'I noticed that you are systematically given the opportunity to state what's gone right and particularly what's gone wrong, and it's important to get these things out on the table and then you find that — at first you hesitate: should I bring this up? Well, why not! — and then you find that you are not the only one; there are four or five others who agree with you [*psychological safety* and *intentional trust*]. You're not alone, it becomes an important issue, you've got it off your chest and a week later the whole thing is resolved, that's how fast it goes.'

'We developed a team tool, in which we incorporated the values that we had put together as a team — we believe this and this and this to be important — we incorporated those into a team tool and we used that to score on a weekly basis so we could measure our progress in time and we could use it to lock on to the things we thought were going well: so when at a given point there was a good deal of friction about something, we could say, well that's a question of project management or the course we were taking or whether or not there was sufficient communication concerning vision or shared goals, you name it, and that led to our being able to resolve these problems within a week [*shared norms*].'

When we compare this group (first column in Table 4) with our reconstruction of a community of practice (Table 3), we see important similarities. The participants were involved in learning during practice. *Tacit* knowledge especially was learned *collectively*, as the example of the application of a certain methodology indicates. The project was described as a joint enterprise. There was *frequent*, *open* interaction, accompanied by both *competence* and *intentional trust*, on the basis of *shared norms*, with indwelling and the building of *empathy*, and perhaps *identification*. This was supported by explicit attention to the team process. This helped to use *trust* as a basis for governance, limiting relational risk, especially *psychological risk*. Thus, risk and *spillover risk* are also reduced by a certain *closure* of the team. However, there are also distances, compared to our rendering of the community of practice. The project team has less *durable ties*: after a specific project is finished, the team is disbanded. It was not ongoing mutual dependence, but a common assignment to a client that created the coherence of the team. As a result, *closure* is less and *cognitive distance* is not so reduced as would typically be the case in a community of practice (if our reconstruction of it is correct). As a result, trust may be less 'thick' than in a community of practice, but this type of group has the advantage of more change in its group composition, yielding the benefits of greater variety and change for innovation.

Expert Group

Another learning group was what we call the 'expert group' (second column in Table 4). Expert groups originate not in a specific common context of application, but in a common content of *generic knowledge*. Participation does not follow from common projects. There was no other common goal other than exchanging *codified knowledge*. Participants met *frequently* around a certain theme (content) in which the participating consultants were especially interested. As far as their daily work is concerned, the participants were involved in very different projects, and some members even came from other divisions within X. The non-redundant contacts (Burt 1992), with *structural holes*, expanded even to the external environment of X, for example, speakers from universities were invited (limited *group stability*). Some of these expert groups were more institutionalized than others. Some obtained a budget and a target (for example, the development of a tool) from the management of the business unit (*some centrality*). Other expert groups were formed spontaneously (*low centrality*). Everybody could become member of an expert group without fulfilling any conditions (*no closure*).

Participants did not expect the knowledge acquired to be directly related to daily work practices. There was not much discussion about the meaning of the acquired knowledge for the individual participant during the meeting (limited *openness of communication*). Hence, the purpose is not so much more efficient exploitation, but *exploration* of novelty in certain specific professional areas. Every participant made sense of the acquired knowledge on his or her own (*individual learning*, *no indwelling*, and *no identification*). This yields limited potential *psychological risk* in loss of reputation or face, because there was no need to give sensitive information about practice that could exhibit weaknesses. One could contribute only information that would enhance reputation. As a result, there was also no *career risk*. There does seem to be a *legitimation risk* in brilliant contributions that might be seen as 'showing off' and putting others to shame, which might jeopardize one's career by being excluded from joint projects.

There does seem to be a high potential *spillover risk*. Being highly codified, the knowledge easily spreads. Participants had high mutual *absorptive capacity* in the specific themes selected per meeting, on the basis of which consultants decided whether to participate. In other words, there was mutual absorptive capacity on the basis of self-selection. This contributes to spillover risk, which is further enhanced by the openness of the group and its instability, with frequent new entrants. No apparent governance was in place to mitigate this risk. Based on our analysis of the conditions of spillover risk, in the theoretical section of this article, our explanation is as follows. In this group, we are dealing with professionals, who operate mostly on *competence trust*, with considerable *openness of communication* concerning only highly codified knowledge. It does not include the sharing of tacit knowledge on how to apply such knowledge in specific projects, or knowledge on customers. While there was a high potential risk of spillover, participants were not afraid of it, since professional knowledge does not yet entail an implementation ability in projects, which is not shared. In view of this, the gain from spillover

seems to be higher than the potential loss. So, there was limited need for building intentional trust. Since there was no need for specific investments, neither for mutual understanding nor for building trust, there was no risk of *lock-in*, and no need for durable ties to recoup specific investments.

Professional Development Group

A third type of group discussed *generic* themes of professional development, such as the norms and values of a consultant, conflict resolution, and giving unsolicited advice to a client. This is related to the aspect of 'being' in learning. Members could not often achieve a full treatment of or final answer to a question during the meeting, and there was no pressure and no attempt to reach consensus either on the general issue or on the personal experience that was brought in. The individual member decided whether and how he utilized the comments. In short, learning was *individual*.

There was much interaction and feedback. This enabled the group to exchange more *tacit knowledge*. They used their own experience in practice as material to work with. However, such experience is abstracted from its specific context, and becomes a hypothetical case: some details are highlighted, others are left out. Context-specific detail need not be fully incorporated. Situations may be discussed that never actually happened. Members can hide in vague references such as 'imagine that' and 'as I once experienced'. Thus experience is turned into 'near histories' and 'hypothetical histories' as critical events for thought experiment (March et al. 1996: 4). Idealizations yield less *reputational risk*. In other words, it is safer to discuss such idealizations, since no one knows to what extent the idealizations or near histories reflect actual weaknesses or failures. This eliminates *psychological risk* and *career risk*. *Spillover risk* is low, since the knowledge involved is general, and has little connection with the specialized competencies by which consultants distinguish themselves. The whole group decided whether a new member was accepted (fairly high *closure*). This closure further reduces psychological and spillover risk. In their daily work, the members were involved in quite different projects. The members reported a symmetrical relationship (*no centrality*). The members of the group might or might not work together in a project. This yields fairly high *cognitive distance*. They met on a regular basis, mostly once a month (*regularly, but not frequently*). Absenteeism is very low, and members were concerned when somebody was absent (*stable* membership). Members used the group as a working space and not just to narrate their experiences. Through discussion, feedback, and critical questions by colleagues, members were invited to explicate some of their implicit assumptions (*open communication* and *competence trust*). Members reported that they learned a lot because their colleagues asked very critical questions that made them reflect. Often, an individual process of reflection was initiated by a colleague's questions. There was intensive feedback, with a short cycle, during the meeting (*indwelling*, building of *empathy*, and perhaps also *identification*). Participants were eager to give comments and feedback to colleagues. In sum, there was high *intentional* trust.

To achieve mutual understanding, some *specific investment* was needed, but it was not high, since the subjects discussed were not specialized. Investment was specific in the sense that people had to adjust to each other's personal experiences and perceptions, but not to the specific contexts of specific projects. The idealizations are applicable more generally, so that, largely, the investment is not specific to the group. However, the building of trust, based on empathy and perhaps identification, did entail a specific investment. Hence, there was a need for a certain *durability* of ties. The closure of the group helped in not disturbing established trust relations. Ties were *intensive*, in the sense that they required a high commitment, but this was a commitment in psychological rather than financial resources.

From Project Improvement to Orientation

A fourth group was intended for project improvement, but for a variety of projects, taking the various projects of group members as input for group discussions. The intention was that a member could present (parts of) a current project and the group could help to improve this project by giving advice. Hence, the scope of learning was *specific*. Often, the agenda was determined by a managing consultant (*centrality*). The meetings took place once a month (*infrequently*). Absenteeism was high and irregular, yielding *instability* of the group. Absence was mostly explained in terms of the priority of the current project or a client, which reflected little commitment, in ties that were *not intensive*. Other group members did not question the reasons for absence. Generally, the meeting ended when every member had had a turn speaking about the current project. Sometimes, the story about a project was interrupted by a question related to the project (and not related to the craftsmanship of the consultant as in the development group). However, there was no pressure from other group members to explain the detailed background and course of a current project. As a member explains:

> 'This would cost too much time, and in this case not every member would get his turn ... It's better if you know more because then you hear more in what is said and it's more interesting; then real dialogue is possible, although it is difficult to do it with the whole group.'

This yields a good illustration of the need to make *specific investments* in mutual *absorptive capacity*. However, the diversity of projects was too high, group membership was too instable, and ties were insufficiently durable to make that worthwhile. Project improvement could hardly take place, as the group did not succeed in discussing projects in sufficient detail. As a result, members reported that they did not learn a lot during the meetings.

A second problem lay in *psychological risk*, in loss of reputation and face and in the ensuing *career risk*, which was intensified by the presence of managing consultants. In contrast with the 'idealizations' in the development group, here, for the improvement of specific projects, the stories would have to be more or less realistic accounts of real events. Members tried to give a positive (or at least not too negative) story. The impression that someone did not function properly in a project had to be avoided.

'There are a lot of ambitions in the group and the participants want to achieve them. They don't want to bring up matters that will hinder them from reaching their goals ... Information that comes out in the group has to be safe. Some things could play a role in one's career. That's why individuals do not want to admit they are incapable of certain things. There are no safety-checks.'

There was also considerable *risk of spillover*. Giving details on projects, including customer contacts, might yield competition. This was intensified by the fact that participants, coming from different projects, spanned many outside *structural holes*. There was a high probability that members of the group would work with outsiders on future projects.

For the group to achieve its purpose, the combination of psychological risk, career risk, and spillover risk should have been mitigated on the basis of trust, and perhaps some closure of the group. However, the group was too unstable, open, and insufficiently dense, and ties were insufficiently durable, to build trust, an internal reputation mechanism, and social control within the group. Also, there was little *mutual dependence*: consultants did not really need each other, in the sense that they did not have a common assignment and there was no pressure to come up with good results together, in a shared project. As a result, there was also little motive to build trust and commitment.

A third reason for lack of learning was that the themes were very time sensitive. When a consultant felt the need for a colleague's advice, this was mostly not the time when the infrequent meetings were scheduled.

As the members admit, for all these reasons willingness to share knowledge and experience was low. The participants were thinking about stopping this type of meeting, if the situation did not change. Summing up, this type of group collapsed due to internal contradictions between its purpose and the conditions needed for achieving it.

However, in spite of the failure to achieve its espoused purpose, members did not wish to abolish the group totally. They like meeting their colleagues and hearing about the projects everybody is currently involved in. This developed into a new enacted purpose. While the group failed in its purpose of project improvement, it served a function of *orientation*, that is, to yield location knowledge on who was involved in what kind of project or had what kind of competence, to be followed up in other contacts outside the group. Contacts are continued in other settings. When an individual member gets the impression that somebody else can help him or her, this person will be contacted outside the group. We would predict that this would then lead to something like the project team discussed before, in which the purpose and features of the group are consistent.

The structural and other features of the group did not obstruct the purpose of seeking only location knowledge. For that, there is no need to provide knowledge that is sensitive with respect to reputation, face, career prospects, or spillover. It can be used to advertise in what area one works, and thus invite referrals for new forms of collaboration. There is no need to engage in specific investments for mutual absorptive capacity (everyone is able to absorb simple location knowledge). Lack of relational risk eliminates the need to build trust in another specific investment. Lack of specific investment prevents lock-in

risk, and enables instability of membership and ad hoc contacts without durable ties. The lack of density and openness of the group, as well as weak ties, is conducive rather than obstructive to the new purpose. A comparison between columns four and five of Table 4 shows that hardly any features of the group had to change in the shift of purpose. All that happens is that managing consultants are no longer present, which eliminates all hierarchy and centrality. The shift of purpose eliminates all relational risk.

Conclusions

We proposed a set of variables for the relevant features of learning groups, concerning knowledge and learning, structure and type of ties, relational risk and its governance (Table 1). We used that to clarify and specify in more detail the notion of 'community of practice' (Table 3), and to reconstruct the four learning groups we found in a consultancy firm (Table 4). We suggest that our framework is useful for a systematic reconstruction of the learning groups we found, and for capturing important distances between them. The hypotheses from the framework (Table 2) yielded an explanation of what happened in the learning groups we found, in terms of consistency between the purpose of the group and features of knowledge and learning, relational risks, and ways to govern them. This helped to explain how the different groups operate and achieve their purpose, or fail to do so (the project-improvement group). The framework was corroborated in the sense that the groups we found did not show the hypotheses to be false. They either confirm them or are 'neutral' with respect to some hypotheses. By 'neutrality' we mean that a hypothesis did not apply to the case, so that it was neither confirmed nor refuted. The study yielded additions and refinements to the framework rather than contradicting it.

As expected, we found no control by formal agreements, very little control by hierarchy, and no control by direct financial incentives. As expected, trust often plays an important role in reducing relational risk.

The analysis of the project team shows that improvement of specific projects, with a high degree of tacit knowledge, and intensive interaction, on the basis of realistic stories of practice, require open communication, yield high potential psychological, career and spillover risk, and require specific investments for mutual understanding, which yields a certain amount of lock-in. Risks are governed by high mutual dependence, internal reputation mechanisms and social control, and the building of both competence and intentional trust, which also entails a specific investment. Potential relational risk and specific investments require a dense and more or less stable structure, and perhaps a certain closure of the group, with strong (intensive, frequent, and durable) ties. This group is close to our interpretation of communities of practice. According to one of our hypotheses (Table 2, hypothesis 7), the problem with this type of group is that it may become so closed, dense, stable, and durable as to reduce cognitive distance and constrain radical innovation. However, compared with our interpretation of communities of practice, the

project team mitigated this problem with greater variety of membership, less closure, and less durable ties. This case is in agreement with all the hypotheses (Table 2), except the last one (hypothesis 7), which is neither confirmed nor contradicted. To test that hypothesis, we would need to test whether in such a group the scope for exploration in order to achieve radical innovation is limited, as we predict.

The analysis of the expert group shows a case in which knowledge is highly codified. It shows that for sharing codified, expert knowledge, with the primary aim of exploration, one needs open groups with variable membership. Because of the choice of specific subjects of expertise, according to which participants decide on attendance, there is a high mutual absorptive capacity. Since one does not need to show weakness, in detailed, realistic accounts of practice, and one can hold back sensitive information, there is no psychological or career risk. There is high potential spillover risk, but this concerns only highly codified knowledge, and not the complementary practice-based tacit knowledge that is needed to compete effectively and quickly. By the time absorbed knowledge is used to compete, it is likely to have changed. In the absence of relational risk, there is no need to build intentional trust. In the absence of the need for specific investments, there is no risk of lock-in, and ties can be short and ad hoc. This case is either consistent or neutral with respect all hypotheses. The argument concerning why apparent spillover risk does not, in fact, apply is not included in the summary of hypotheses in Table 2, but in the more detailed discussion in the theory section of this article.

The analysis of the project-improvement group shows an interesting, illustrative failure. This group collapsed from inconsistency of purpose and features of the group. Its purpose was mainly oriented toward exploitation, while its structure favoured exploration. To achieve its purpose, the group would need more resemblance to the project team. Its purpose required the sharing of detailed, tacit, practice-based project knowledge, which would require specific investments in mutual absorptive capacity. That was frustrated by the diversity of projects and the limited duration of ties. In view of the task, there were high psychological, career, and spillover risks, which were intensified by the openness of the group, by the presence of managing consultants, and by participants bridging many external structural holes. Relational risk required the building of trust, internal reputation, and social control, but these were blocked by the lack of density, stability, intensity, frequency, and duration of ties. Interestingly, a shift of purpose from project improvement to the sharing of only location knowledge enabled the group to function well: all its features were conducive to the new purpose, which required none of the conditions that the group lacked. This case is consistent with all hypotheses except hypothesis 7 (which is neutral). The hypotheses predict that the group could not but fail in its original purpose.

The analysis of the professional development group shows an interesting, novel insight that adds to our theoretical framework. Here, potential relational risk was eliminated partly according to hypotheses 5 and 6, but also partly by a new device. This was the use not of detailed, realistic, project-specific stories, which were not needed here, but of ‘idealized’, ‘hypothetical’ ‘near

histories' of practice. In terms of our framework, this serves several purposes. First, it yields a reduced need for specific investment in mutual absorptive capacity. There is no need for lengthy explanation of all requisite detail, replete with tacit, context-specific knowledge. This reduces the lock-in risk. Second, it greatly reduces psychological, career, and spillover risks. This allows for open communication. It also allows for less stable ties, with limited duration, cognitive distance, and some openness of the group, with crossing of external structural holes. This enhances the potential for exploration.

Other lessons from the study yield the following additions and refinements to the theoretical framework.

1 In addition to the stability, duration, and frequency of ties, which are included in the framework, there is a new aspect of regularity (professional development group). This may yield a worthwhile refinement. Next to duration and frequency, regularity serves to enhance the intensity of ties.
2 Time specificity is also a relevant aspect for project improvement: when one needs help, one needs it urgently, and infrequent meetings will not suffice (project-improvement group).
3 Spillover risk includes not only competence-related knowledge, but also commercial knowledge about customers (project-improvement group).

We see the following as priorities for further research.

1 More extensive, statistical testing of the framework, on the basis of larger samples.
2 A test of hypothesis 7, which states that groups such as the project team and communities of practice restrict the scope for radical innovation.
3 More attention with regard to the position of individuals within groups (centrality and boundary spanners) and within-group dynamics (indwelling and resolution of conflict).
4 Relational risk in groups might be lessened or amplified by organizational culture. This should be taken into account in future research. In particular, one might compare groups that are structurally similar, but arise in different organizations with different cultures in order to identify possible culture effects. For a detailed empirical study of the interpersonal effects of conflict on the deepening or breaking of trust, and the effect of organizational culture, refer to Six (2002).
5 We used the theoretical framework for an *ex post* explanation of learning groups that were identified empirically. The framework might also be used to predict viable forms, and then test their occurrence.

References

Archer, Margaret S.
1995 *Realist social theory: The morphogenetic approach*. Cambridge: Cambridge University Press.

Argyris, Chris, and Donald Schön
1978 *Organizational learning*. Reading, MA: Addison Wesley.

Bateson, Gregory
1973 *Steps to an ecology of mind*. London: Paladin.

Belbin, Meredith
1984 *Management teams: Why do they succeed or fail?* London: Butterworth.

Blackler, F.
1995 'Knowledge, knowledge work and organizations: An overview and interpretation'. *Organization Studies* 16/6: 1021–1046.

Brown, John S., and Paul Duguid
1996 'Organizational learning and communities of practice' in *Organizational learning*. M. D. Cohen and L. S. Sproull (eds), 58–82. London: Sage.

Bruner, Jerome
1979 *On knowing: Essays for the left*
[1962] *hand*. Cambridge, MA: Belknap Press.

Burt, Ronald
1992 *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.

Cohen, M. D., and P. Bacdayan
1996 'Organizational routines are stored as procedural memory' in *Organizational learning*. M. D. Cohen and L. S. Sproull (eds), 403–430. London: Sage.

Coleman, J.
1990 *Foundations of social theory*. Cambridge: Cambridge University Press.

Cook, S. D. D., and D. Yanow
1996 'Culture and organizational learning' in *Organizational learning*. M. D. Cohen and L. S. Sproull (eds), 430–450. London: Sage.

Edmonson, Amy
1999 'Psychological safety and learning behaviour in work teams'. *Administrative Science Quarterly* 44: 350–383.

Fiol, C. M., and M. A. Lyles
1985 'Organizational learning'. *Academy of Management Review* 10/4: 803–813.

Frey, Bruno S.
2002 'What can economists learn from happiness research? *Journal of Economic Literature* 40/2: 402–435.

Giddens, Anthony
1984 *The constitution of society*. Cambridge: Polity Press.

Granovetter, Mark
1982 'The strength of weak ties' in *Social structure and network analysis*. P. Marsden and N. Lin (eds), 105–130. Beverly Hills, CA: Sage.

Granovetter, Mark
1992 'Problems of explanation in economic sociology' in *Networks and organizations*. N. Nohria and R. G. Eccles (eds), 25–56. Boston, MA: Harvard Business School Press.

Hansen, M. T.
1999 'The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits'. *Administrative Science Quarterly* 44: 82–111.

Hedberg, B. L. T., P. C. Nystrom, and W. H. Starbuck
1976 'Camping on seesaws: Prescriptions for a self-designing organization'. *Administrative Science Quarterly* 21: 41–65.

Hirschman, Albert O.
1970 *Exit, voice and loyalty: Responses to decline in firms, organisations and states*. Cambridge, MA: Harvard University Press.

Janis, Irving
1981 *Groupthink*. Boston, MA: Houghton Mifflin.

Krackhardt, David
1999 'The ties that torture: Simmelian tie analysis in organizations'. *Research in the Sociology of Organizations* 16: 183–210.

Lave, Jean, and Etienne Wenger
1991 *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.

Lewicki, R. J., and B. B. Bunker
1996 'Developing and maintaining trust in work relationships' in *Trust in organizations: Frontiers of theory and research*. R. M. Kramer and T. R. Tyler (eds), 114–139. Thousand Oaks, CA: Sage.

McAllister, D. J.
1995 'Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations'. *Academy of Management Journal* 38/1: 24–59.

Maguire, S., N. Philips, and C. Hardy
2001 'When "silence = death", keep talking: Trust, control and the discursive construction of identity in the Canadian HIV/AIDS treatment domain'. *Organization Studies* 22/2: 285–310.

March, J. G., L. S. Sproull, and M. Tamuz
1996 'Learning from samples of one or fewer', in *Organizational learning*. M. D. Cohen and L. S. Sproull (eds), 281–288. Thousand Oaks, CA: Sage.

Merleau-Ponty, M.
1964 *Le visible et l'invisible*. Paris: Gallimard.

Nooteboom, Bart
1999 *Inter-firm alliances: Analysis and design*. London: Routledge.

Nooteboom, Bart
2000 *Learning and innovation in organizations and economies*. Oxford: Oxford University Press.

Nooteboom, Bart
2002 *Trust: Forms, foundations, functions, failures and figures*. Cheltenham: Edward Elgar.

Nussbaum, M. C.
2001 *Upheavals of thought: The intelligence of emotions*. Cambridge: Cambridge University Press.

Piaget, Jean
1970 *Psychologie et epistémologie*. Paris: Denoël.

Piaget, Jean
1974 *Introduction a l'épistémologie génétique*. Paris: PUF.

Polanyi, Michael
1962 *Personal knowledge*. London: Routledge.

Simon, Herbert A.
1983 *Reason in human affairs*. Oxford: Blackwell.

Six, Frederique
2002 'Trust and trouble in organizations' in *The trust process in organizations*. B. Nooteboom and F. Six (eds). Cheltenham: Edward Elgar.

Von Krogh, Georg
1998 'Care in knowledge creation'. *California Management Review* 40/3: 133–153.

Vygotsky, Lef S.
1962 *Thought and language*, ed. and trans E. Hanfmann and G. Varkar. Cambridge, MA: MIT Press.

Wegner, D. M., T. Giuliano, and P. Hertel
1985 'Cognitive interdependence in close relationships' in *Compatible and Incompatible Relationships*. W. Ickes (ed.), 253–276. New York: Springer.

Weick, Karl F.
1979 *The social psychology of organizing*. Reading, MA: Addison Wesley.

Weick, Karl F.
1995 *Sensemaking in organizations*. Thousand Oaks, CA: Sage.

Weick, Karl F., and Frances Westley
1996 'Organizational learning: Affirming an oxymoron' in *Handbook of organization studies*. S. R. Clegg and W. R. Nord (eds), 440–458. London: Sage.

Wenger, Etienne C., and William M. Snyder
2000 'Communities of practice: The organizational frontier' *Harvard Business Review*, January–February: 139–145.

Zand, D. E.
1972 'Trust and managerial problem solving'. *Administrative Science Quarterly* 17/2: 229–239.

Irma Bogenrieder

Irma Bogenrieder is Assistant Professor at the Rotterdam School of Management. Her interests lie in the field of organizational behaviour, especially in the functioning of groups within an organization. Currently, she is focusing on learning within diverse types of groups, concentrating on both the dynamics within the group and conditions for the formation and persistence of the group.
Address: Rotterdam School of Management, Erasmus University Rotterdam, PO Box 1738, 3000 DR Rotterdam, The Netherlands.
E-mail: ibogenrieder@fbk.eur.nl

Bart Nooteboom

Bart Nooteboom is Professor of Organizational Dynamics at the Rotterdam School of Management. His interests lie within the broad field of entrepreneurship, innovation, and learning between and within organizations. He attempts to combine perspectives of competence in types of knowledge and learning, and governance for the management of 'relational risk'. He is member of the Netherlands Academy of Sciences (KNAW). He has published widely in such journals as *Organization Studies*, *Academy of Management Journal*, *Journal of Management and Governance*, *Cambridge Journal of Economics*, *Journal of Evolutionary Economics*, *Journal of Economic Dynamics and Control*, *American Journal of Sociology*, *Research in the Sociology of Organizations*, *Journal of Computational and Mathematical Organization Theory*, *Small Business Economics*, and *Research Policy*. His published books include *Inter-Firm Alliances: Analysis and Design*, *Learning and Innovation in Organizations and Economies*, and *Trust: Forms, Foundations, Functions, Failures and Figures*.
Address: Rotterdam School of Management, Erasmus University Rotterdam, Burgemeester Oudlaan 50, PO Box 1738, 3000 DR Rotterdam, The Netherlands.
E-mail: bnooteboom@fbk.eur.nl