

Doing Business with Theory: Communities of Practice in Knowledge Management

Norman Makoto Su¹, Hiroko N. Wilensky² & David F. Redmiles²

¹*School of Information and Library Studies, University College Dublin, Dublin 4, Ireland (E-mail: normsu@ucd.ie);* ²*Department of Informatics, University of California, Irvine, Irvine, CA 92697-3440, USA (E-mail: hwilensk@ics.uci.edu; E-mail: redmiles@ics.uci.edu)*

Abstract. We explore how the notion of *communities of practice* (CoPs) was translated and popularized from its original inception by Lave and Wenger in 1991. We argue that the Institute for Research on Learning (IRL), a spin-off of Xerox PARC, proved instrumental in enrolling CoPs into the knowledge management (KM) discipline. IRL objectified, packaged, and made a business out of CoPs. CoPs in KM are now a formalized process coupled with technological artifacts to build groups of people who effectively share knowledge across boundaries. Drawing from participant observations, archival documents, and interviews with KM practitioners in the aerospace industry as well as key players of IRL, our research seeks to unveil the invisible history that the popularization of a theory can often obscure. We argue that CoPs provide a case study for understanding how abstract concepts in science are strategically and subconsciously reified, or made objects of inquiry, and appropriated by actors. This reification of a “soft” science blurs the line between theory and technology.

Key words: aerospace, communities of practice, knowledge management, science & technology studies, sociology of scientific knowledge

1. Introduction

Walk into your local bookstore and amble over to the business or management section. Leaf through any book on knowledge management (KM), and you are likely to see a chapter or two on *communities of practice* (CoPs). Rumizen’s (2002) popular book—*A Complete Idiot’s Guide to Knowledge Management*—is a good case in point. Below, we quote her introduction to Chapter 8 (p. 85), entitled “Communities of Practice—The Killer Application”.

In This Chapter

- Characteristics of communities of practice
- Role of the community coordinator

– Launching a community of practice at SAP America

I dare you to find a serious general conference on knowledge management that doesn't include at least one session on communities of practice. Slowly over the past few years, communities of practice have come to be acknowledged as the *killer application for knowledge management*. And rightfully so. [emphasis added]

Theories are a “dime a dozen,” as the cliché states. However, CoPs are special. When Jean Lave and Etienne Wenger (1991) first wrote about communities of practice, it was part of a conceptual framework, which included legitimate peripheral participation, meant to provide a new analytic lens into learning. Far from being merely “an informal label for a knot of ideas” (Lave 2008, p. 283) or an “intuitive notion” (Lave and Wenger 1991, p. 42), CoP has become a hugely influential, concrete social theory in academia. It has also entered the daily discourse of those whose lives intersect management. More specifically, it has become part and parcel of a particular discipline of management called knowledge management. Those who claim to be knowledge management practitioners or professionals must eat, drink, and breathe CoPs.

The phrase ivory tower is often employed to describe how academics are out of touch with the “real” world, but we believe that many researchers would welcome a chance for their theory, methodology, or invention to become relevant to the everyday lives of people. For some disciplines, how such things can become useful to the layperson is apparent. Yet, rare is the case for the more “abstract” results of social studies to escape the academic hallways. This makes the case for CoP all the more impressive.

Scholars have long noted the polysemous nature of CoP. Rather than reinvent the wheel, this paper focuses on the *business* of communities of practice. More specifically, we are interested in the processes by which actors were convinced, recruited, and then aligned to ensure that CoP became an indispensable actor in the knowledge management discipline. We focus not on the changing discourse of CoP in academia, but rather on how CoP was translated to the masses. Simply put, how did an academic theory become so popular?

Like many ethnographies, from the outset our inquiry into CoP was planned not deliberately but circuitously through our investigation of the KM profession. During our observations, we noted how deeply ingrained CoP was in the KM discourse. Moreover, how these KM practitioners understood CoP was drastically different from how we, as academics, had learned about CoP. Here, in the wild, we had spotted an analytic perspective supposedly promulgated through established methods of scientific communication. This spurred us to go back to the origins, when CoP was born. Might a historical perspective on CoP hint at how it achieved its popularity?

In this paper, we present a partial history of how CoP has become an integral actor in KM. First, drawing from our ethnography of KM practitioners

Doing Business with Theory: Communities of Practice in KM

in the aerospace industry, we ask how CoP is currently conceived by practitioners. Second, we go back and investigate CoP's inception at the Institute for Research on Learning (IRL). There is remarkably little information in the "history books" on this offspring of the Xerox Palo Alto Research Center (PARC). Although CoP may be an idea turned theory, it nonetheless became a product—one that was packaged, marketed, shipped, and sold to customers. Treated as one might treat any tangible object, CoP presents a fascinating case study of the popularization of social theories. It is our contention that CoP has morphed from its original form and has become objectified, and further, that objectification has proven useful when conducting business about CoP. Finally, we argue that this elevation of perspective to theory blurs the line between science and technology for the "soft" sciences.

2. Critical perspectives on communities of practice

It is difficult to pinpoint one definition of communities of practice. Indeed, past work has reflected upon CoP as manifested in its seminal works:

1. *Situated Learning: Legitimate Peripheral Participation* by Lave and Wenger (1991)
2. *Organizational Learning and Communities-of-Practice: Toward a Unified View of Working, Learning, and Innovation* by Brown and Duguid (1991)
3. *Communities of Practice: Learning, Meaning, and Identity* by Wenger (1998)

Rather than attempt to re-summarize what CoPs are (in their various contexts), we refer the reader to two excellent reviews by Cox (2005) and Osterlund and Carlile (2005). Other reviews exist that focus on particular aspects of CoP (Contu and Willmott 2003; Davenport 1994). Nonetheless, we believe it constructive to discuss here when these scholars go noticeably beyond a neutral stance on CoP's various instantiations.

Cox (2005) provides a comparative review of how CoP has evolved over time in the three cited. He provides a more general overview of how CoP's conceived utility has changed (p. 527):

Sometimes it is a conceptual lens through which to examine the situated social construction of meaning. At other times it is used to refer to a virtual community or informal group sponsored by an organisation to facilitate knowledge sharing or learning.

For Cox, CoPs in their various instantiations differ across their conceptualization of community, learning, power, change, formality, and diversity. Osterlund and Carlile (2005) use a relational framework of practice theory to parse out the different guises of CoPs. Their framework contains seven attributes of relationship thinking that share some overlap with Cox's own method of analysis; for

example, it asks, “Where does the author introduce the notion of power?” and “What are the dependencies associated with a relation?”

In its first appearance (Lave and Wenger 1991), community of practice was never explicitly defined; it is only in later works has the concept become a more explicit *object* of concern. Cox notes weaknesses in the analytic lens proposed by Lave and Wenger (1991): (1) CoP is “not adequate to explain all the power forces within a community, let alone those that structure it from outside”; and (2) “[T]he relationship between communities or between communities and other entities as a source of change and conflict is not considered.” Osterlund and Carlile (2005) agree with the latter point, noting that “[t]he danger...in this particular framework...is that it depicts communities of practice as largely independent and unconnected.”

The highly cited article by Brown and Duguid (1991) firmly places CoP into a workplace context by melding Orr’s (1996) ethnography of Xerox photocopier repairmen with Lave and Wenger’s own analytic approach to learning. One point of contention brought up by both reviews is the undue emphasis on the homogeneous aspects of communities. Cox calls their notion of CoP as being “internally egalitarian” and “losing sight of conflict.” Osterlund and Carlile (2005) also include a follow-up article— “Knowledge and Organization: A Social-Practice Perspective” (Brown and Duguid 2001)—when critiquing CoP’s harmonious image: “it is difficult to find much reference to *power* relations when it comes to how Orr’s service technicians interact...[T]hey end up reproducing a notion of community...depicting a coherent group of people organized around a set of shared characteristics.”

Indeed, Cox (2005) and Osterlund and Carlile (2005) take issue with the data set on which Brown and Duguid ground their theoretical constructs. One issue that Osterlund and Carlile note is that the technicians case study is inadequate in understanding multiple communities: “[w]e are left craving good empirical examples describing cross-communal relations...[They] hinge on weak empirical material.” A second issue that Cox suggests is that the empirical data may not be representative: “[t]here are good grounds for seeing the repairmen’s situation as rather rare...[It] can hardly be regarded as typical work conditions; vary the conditions, and the resulting counter-community, if one emerged at all, might be very different.”

In analyzing Wenger’s (1998) magnum opus, Cox (2005) states that identity as maintained through “multi-membership of different communities” becomes the narrative of CoP. This multicommunal perspective, as Osterlund and Carlile call it, potentially addresses one of the deficiencies in Lave and Wenger’s (1991) original report. However, Wenger’s (1998) data still draws from one population (insurance claim processors) and retains the homogeneous flavor of Brown and Duguid’s work: “the notion of identity...solely defines the similarities among individuals.” Cox especially considers the term “community” as problematic: “it almost becomes difficult to see why Wenger used the term ‘community’ at all

Doing Business with Theory: Communities of Practice in KM

since he denies most of our usual assumptions about it...as a rather large, helpful and friendly, bounded group.”

Finally, and perhaps most relevant to our paper are the discussions (Cox 2005; Osterlund and Carlile 2005) on the relationship between management and communities of practice. Cox believes Wenger (1998) “underestimates the powerful rationalizing processes in capitalism and the ability to rapidly appropriate and systematize understanding.” Additionally, Cox asserts that despite the rhetoric of CoP as being radical, subsequent developments, such as the supposed knowledge sharing tool Eureka (to Cox, a simple database rather than a storytelling tool), actually only reinforces that CoP “can be channeled through rather familiar rationalizing processes” by management.

In his review, Cox also includes Wenger et al.’s (2002) book as the site where “commodification of the idea of community of practice” occurs. Cox considers the concept of CoP here as radically different: “the idea is to create or foster new groupings of people...to invent new practices.” Again, as mentioned in the previous paragraph, there is no discussion of the dangerous potential for CoP to be leveraged by management to control its employees and to simply recreate another new norm to “impose on participants, oppressive in the same way the ‘team’ can be.” Management is now part of defining which CoPs are worthy of identification and support. This would go against the counter-cultural or revolutionary rhetoric of Brown and Duguid, who “take the perspective of the service technicians and how they feel squeezed by a community of managers who do not seem to comprehend their...improvisational practices” (Osterlund and Carlile 2005). We see here that the leap from academia to management is fraught with consequences.

Certainly, the appropriation of CoPs by KM practitioners seems to result in the most strongly worded criticisms by these two works (Cox 2005; Osterlund and Carlile 2005). For example, Osterlund and Carlile state that “the three seminal works often gets distorted when adopted by other scholars or practitioners.” In KM, CoPs are seen as “the solution to any knowledge-sharing problem faced by an organization.” Cox concurs, saying, “communities of practice are also claimed to offer solutions to classic management issues.” In fact, CoPs neatly encapsulate what makes KM unique: “community of practice is the classic conceptualization of knowledge management as more than information management: a social not individual or technological solution, about tacit not codified knowledge.”

To summarize, we see how critical views of CoPs have aimed at power and tension, multicommunal relationships, the data upon which it has been founded, and management’s role. There is a definite consensus that CoPs have transformed over the course of these seminal works. The goals of Cox (2005) and Osterlund and Carlile (2005) are not to explain why or how these disparities arose, but rather to state what these disparities are. Having said that, Cox contemplates, “[i]f an explanation for these divergences is sought, the weakness (or perhaps the

strength) of the concept of communities of practice lies in the ambiguity of both the terms ‘community’ and ‘practice.’” We agree that indeed the lexical properties of CoPs have proven vital to its success and survival across varied disciplines, but this is only part of the puzzle. It is precisely our motivation in this paper to come closer to understanding how CoPs came to be. The two cited reviews drew from sources that are considered the primary modes of scholarly dissemination: journals and books. Cox, Osterlund, and Carlile have certainly elucidated a key component of how science works, and they primarily tell the story of how an academic theory has evolved through the academy. Our paper seeks to tell a partial story in what allowed CoPs to enter the knowledge management realm—in other words, how CoPs moved beyond the boundaries of academia into a more public discourse.

3. Related work

Vann and Bowker (2001) take an approach similar to ours in focusing on the concept of “practice.” Specifically, they go back and examine Jean Lave’s (1988) *Cognition in Practice*. Informed by actor-network theory, Vann and Bowker frame practice as a scientific concept that goes against the traditional theories of learning and organization. Formerly, practice is seen as an object one can dissect from its context and then re-situate to be taught in any classroom. However, practice as seen by Lave is situated in a social practice constituting a spatial and temporal context that cannot be ignored.

Vann and Bowker (2001) note that ethnography becomes a tool for uncovering the “truth” behind work (Orr 1996). The discourse of consultancy texts and book blurbs demonstrate “how a concept of practice is reinstrumentalized and reconfigured as a commercial object with specific users” (Vann and Bowker 2001, p. 248). Finally, they note that practice is first “discovered” (communities of practice) to already exist, and then it is adopted, cultivated and promoted. Our study examines the business side of CoP and the role that the Institute for Research on Learning had in forming the current conception of CoP prevalent in the aerospace industry.

This work builds upon our ethnographic studies of knowledge management practitioners in the aerospace industry (Su et al. 2007; Wilensky et al. 2008, 2009). These studies focused on the individual beliefs and practices of KM practitioners in their respective corporations as well as in a public forum called the *KM Exchange* (a pseudonym) in which aerospace KM practitioners gathered. This KM Exchange was labeled a community of practice by its constituents.

In the first study, Su et al. (2007) showed that this CoP was primarily a forum for legitimizing the KM discipline itself. Three recurrent themes were found in the KM practitioner discourse: (1) KM makes effective use of knowledge by capturing and reusing it (i.e., knowledge is objectified); (2) KM practitioners’ personal traits and tools/practices are more progressive compared to those of

Doing Business with Theory: Communities of Practice in KM

practitioners in other fields such as IT; and (3) KM is misunderstood within the KM practitioners' work organizations. The second study (Wilensky et al. 2008) found that the asymmetrical power relationships existing in each practitioner's workplace were transferred and mirrored into the KM Exchange. This impacted the boundaries of the community, the knowledge-sharing practices among members, and the institutionalized beliefs within the community. Finally, the third study (Wilensky et al. 2009) explored the challenges KM practitioners faced when disseminating KM ideas in their organizations. For example, practitioners would sometimes blindly adopt preferred solutions by self-labeled KM experts/gurus instead of first observing how such tools might fit the actual work practices of aerospace employees.

In contrast to our earlier work, this paper concentrates squarely on one particular tool of the aerospace KM practitioners' toolkit, communities of practice.

4. Methodology and field sites

Our ethnographic inquiry was conducted in two phases. First, we focused on a community of knowledge management practitioners in the aerospace industry. Second, we investigated the now defunct Institute for Research on Learning. The initial impetus was in understanding the disciplinary discourse of KM practitioners in the aerospace industry. During the course of our fieldwork, however, we were struck by how often KM practitioners frequently and casually discussed communities of practice. Moreover, we noticed that the CoPs discussed at these meetings were significantly different from what we had learned as academics in the informatics field (Dourish et al. 2008). This motivated us to investigate the landscape in which CoP was conceived and how CoP was able to successfully enter the business world.

4.1. The KM Exchange

Phase one lasted for three years and began in 2006 with fieldwork of the KM Exchange (~85 members) in the aerospace industry. The members were primarily KM practitioners from six major aerospace organizations. Others included people from three local universities. Our data consist of (1) 24 semistructured interviews lasting 30~90 min each; (2) observations of ten quarterly meetings, three KM conferences and one seminar; and (3) documents relating to CoP "implementations" in the aerospace industry. Nineteen of the informants were from the aerospace industry and 4 were from academia (with 1 academic interviewed twice). Among the 19 aerospace members, 12 worked for a KM team or department within their work organizations and 7 were affiliated with the KM department/team to some extent. The backgrounds of the aerospace KM members varied, including thermal engineer, rocket scientist, print shop manager, IT

professional, and special librarian. Finally, we gathered archival data, such as PowerPoint slides, handouts, and website text, from the KM Exchange.

4.2. The Institute for Research on Learning

Phase two, motivated by our discovery of the pervasiveness of CoPs in the KM Exchange's discourse, drew from nine semi-structured interviews and archival document reviews to gain a historical perspective on the Institute for Research on Learning. We interviewed former employees of IRL who were key players in the formation of the CoP concept as well as more peripheral researchers who "experienced" the propagation of CoPs. All of our informants worked in IRL sometime between 1986 to 2000. We contacted these individuals for interviews through our professional contacts or directly through email. Historical materials on IRL were digitally scanned from PARC's Information Center (library) archives. Table 1 provides a brief overview of the informants we interviewed.

4.3. Analytic methods

Using the grounded theory approach (Strauss and Corbin 1998), field notes and interviews were transcribed, coded, and analyzed on Atlas/ti (www.atlasti.com) to facilitate the formation of research findings and themes. This approach enabled us to in situ investigate new domains and topics during our iterative data collection and analysis. We have continually scrutinized our data since initial fieldwork in 2006. For this particular article, we revisited and analyzed the data around one category—the KM Exchange members' perceptions of CoPs.

We drew from discourse analysis (Phillips and Hardy 2002) when examining the historical materials of IRL. Discourse is "an interrelated set of texts, and the practices of their production, dissemination, and reception, that brings an object into being" (p. 3). It seeks not *what* the texts are saying but rather *how* the texts are saying what they say in their social settings: "Discourse analysis explores how texts are *made* meaningful through...processes and also how they contribute to the constitution of social reality by *making* meaning" (p. 4).

Overall, our analysis is also implicitly informed by actor-network theory (Callon 1996). One might regard the study of scientific theories as simply the study of science in general (e.g., when scientists empirically discover the scientific laws of the universe). In other words, the biggest ontological discrepancy between technology and theory is that the former seems more "tangible" and more "practical." We believe, however, that methods from actor-network theory prove just as applicable in deconstructing how theories are taken up by actors to make themselves indispensable in their scientific discipline. More to the point, we argue that *the reification of "theory" is a maneuver that facilitates its popularization*. And it is exactly this process of reification, from the

Doing Business with Theory: Communities of Practice in KM

Table 1. Overview of IRL informants.

Name	Year joined IRL	Position prior to joining IRL	Position at IRL/How they got involved in IRL	Current position
John Seely Brown (JSB)	1986	Chief Scientist at Xerox and Director at PARC	Cofounder/Founder of IRL	Visiting scholar and advisor to the Provost at University of Southern California, Independent Co-Chairman of the Deloitte Center for the Edge
Melissa Cefkin	1993	PHD student at Rice University (Anthropology)	Research Scientist/Hired after graduate school	Anthropologist at IBM Almaden Research Center
Paul Duguid	1987	Worked in a New York publishing company	Editor of IRL Reports and Member of the Research Staff/Hired for his publishing experience	Adjunct Professor at University of California, Berkeley (School of Information)
Shelley Goldman	1989	Research Scientist at Center for Children & Technology at Bank Street College of Education	Senior Research Scientist and Director of School & Community Programs/Contact through Roy Pea (Senior Research Scientist at IRL)	Professor at Stanford University (Education) and Director of Learning Design & Technology Program
Brigitte Jordan (Gitti Jordan)	1987	Professor at Michigan State University (Department of Anthropology)	Senior Research Scientist/ Recruited by JSB	Consulting Corporate Anthropologist
Jean Lave	1987	Professor at University of California, Irvine (School of Social Science)	Research Scientist/Invited during sabbatical from UC Irvine	Professor at University of California, Berkeley (Education and Geography)
Charlotte Linde	1989	National Research Counsel Senior Research Associate at NASA Ames Research Center	Senior Research Scientist/Started as a visitor of IRL seminars and eventually hired	Senior Research Scientist at NASA, Ames Research Center
Susan Stucky	1987	Center for the Study of Language & Information at Stanford University	Assistant Director, Senior Research Scientist, and Director of Strategic Practice Group/Recruited by JSB	Manager of the Service Design group at IBM Almaden Research Center
Etienne Wenger	1987	PHD student studying artificial intelligence at University of California, Irvine (Department of Information & Computer Science)	Research Scientist/Recruited by JSB	Independent Researcher, Consultant, Author and Speaker

lens of the sociology of scientific knowledge and actor-network theory, that we wish to investigate.

4.4. Anonymity: caveats of interviewing

We anonymized the data from phase one of our study, but we chose not to anonymize the data from IRL because we felt identity is necessary to tell the story of CoPs in which the historical context is important. Who the people are and how they were situated in IRL are important pieces of information on how the analytic lens of CoPs traversed through various social worlds. Thus, we felt attaching the interviewee's names to their quotes were important and meaningful.

As a result, some of the informants explicitly asked to vet their own quotes, even though we tried to avoid including material that might prove personally sensitive. We agreed to allow these informants to each separately view the quotations we chose to include in the paper. To avoid hearsay, we strove to include only quotations that could be verified by at least one other data source (e.g., publications or another informant). To ensure a rigorous analysis, all coding of transcripts was performed independently of any informant feedback. We believe this achieves a balance between allowing us to tell the deep story of communities of practice while preserving the informants' personal wishes.

Therefore, we issue a disclaimer that the reader is not to interpret the quotations as the final, official opinion of the informants. For many academics, published sources are seen as the authoritative source for theories or commentary. The semistructured interviews cannot be held to a similar standard—and they are not intended to be. The interviews we conducted are a snapshot of a particular informant's view of a particular time. Our primary goal in this paper, however, is to see precisely how people recollect events (in various ways) and is not to construct a precise timeline of events as they transpired. We are not trying to reconstruct a definitive edition of the history of communities of practice in the traditional sense that historians do (Butterfield 1965). Instead, our interest in informants' answers revolves around questions such as: What do they omit? What do they remember? What are their own personal, unique histories of the context surrounding communities of practice? Communities of practice, as we discuss in this article, is in the eye of the beholder.

5. Formalized CoPs in the aerospace industry: the gold standard of communities

Communities of practice is the killer application for knowledge management practitioners (Rumizen 2002). In aerospace companies, CoP are conceptualized by knowledge management teams as a valuable technology for corporations to effectively generate, exchange, and share knowledge. Organizations may have other officially labeled variants of communities, but an overarching theme of KM practitioners places

Doing Business with Theory: Communities of Practice in KM

CoP as the “gold” standard of communities. Within the black box that is the CoP, members are responsible for producing the highest-quality knowledge.

Often, CoP are contrasted with other forms of groups. Previous forms of collaboration are “transformed” or “evolved” into CoPs. For example, the Raytheon Logistics *Council* is transformed into the Raytheon Integrated Logistics *Community of Practice* (RILCOM) (Palla and Ward 2006). Whereas the council has “limited cross business knowledge sharing” and “limited standardization of processes,” the CoP shares “knowledge broadly across organizational and functional boundaries,” revealing “expertise and tacit knowledge faster” (Figure 1). One head of a KM department noted that an Integrated Product Team (IPT) started out as a *Tiger Team* but transformed into a *CoP*. CoPs are placed at the higher end of organizational forms.

In turn, CoPs themselves are set in contrast to other, more informal communities. A chapter in *Knowledge Retention: Strategies and Solutions* written by members of the Aerospace Corporation’s KM group (Sutton et al. 2008) details their corporate strategy integrating KM. Of particular interest is how communities are part of their “KM Roadmap” (p. 94–96):

Addressing cross-organizational, long-term stewardship of technical content in this manner was coordinated through two complementary frameworks: communities of practice (CoPs) and communities of interest (CoIs).

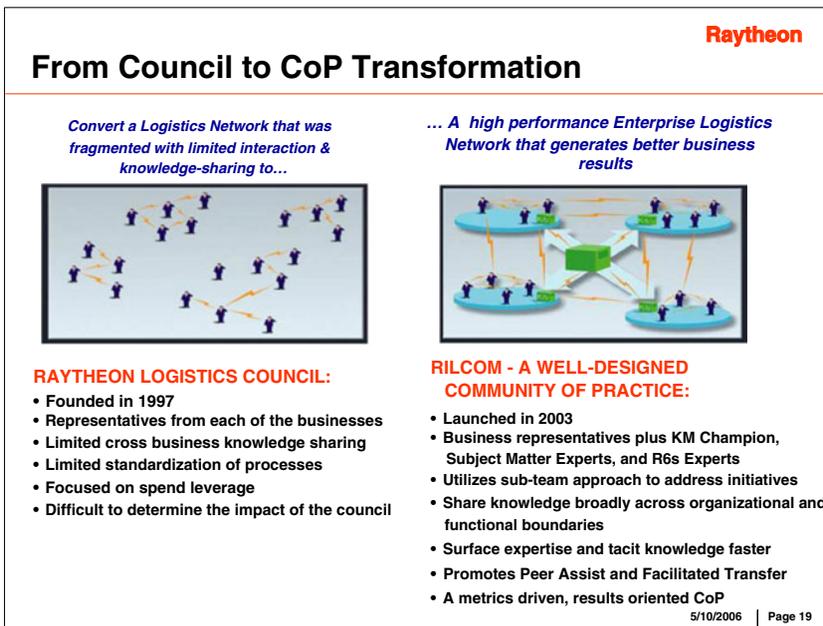


Figure 1. Raytheon: From Council to CoP Transformation (source: Palla and Ward 2006).

For the Aerospace Corporation, these two entities represent opposites in a spectrum of “formality” for communities of practice. A CoI has virtually no formality and oversight. A CoP, however, has “a mandated governance framework of defined roles within the community” (p. 96).

The formal mechanisms for a community of practice allow it to have a single “voice.” Members can work together to publish a community position on its topics. In fact, this formalization gives the community opinion more weight and legitimacy: “this community position has increased value because of the defined process for position formulation, peer review, and approval set forth within the CoP” (Sutton et al. 2008, p. 96–97). This “rigor” allows CoP members to argue that their opinions are well thought out. Those communities with less formalization, such as communities of interest, are relegated to producing knowledge that cannot be representative of a company.

In the next section, we discuss the characteristics that embody the KM practitioners' *formalized* perspective on CoPs in the aerospace industry: (1) CoP as a bridging agent for disparate groups, (2) the prescribed roles in a CoP, (3) the tight coupling between technology and CoP, and (4) the measurement of CoP.

5.1. Bridging disparate groups

Key amongst the aerospace industry's conception of CoP is as a bridging agent for disparate groups. For many KM practitioners, the secretive nature of the aerospace industry poses a serious impediment to the tenets of KM—knowledge sharing (Su et al. 2007). Employees are placed in “silos” or “stovepipes.” In these autonomous groups, there is little motivation or opportunity for cross organizational sharing. According to NASA, one problem KM seeks to solve is “geographic dispersal.” Leonard and Kiron (2002) note that “NASA's workforce is spread across 10 different centers in eight states” and that “this makes communication and collaboration difficult.”

For KM practitioners, CoPs provide a solution for these islands of knowledge. For example, a CoP at the one aerospace corporation is described in a slide as “bridging many training empires.” CoPs provide a unifying element for “multiple sites, companies, [and] teams.” At one aerospace and defense corporation, CoPs provide “horizontal integration” and helps “knowledge flow across organizational and geographical boundaries.” At the Air Force (a primary customer of aerospace companies), knowledge management converts stovepipes into communities through the Air Force Knowledge Now (AFKN) strategy (Figure 2).

One KM leader explained how, through CoPs, she could successfully connect employees in her company who were geographically apart statewide:

Right now we're doing a [Company Name] engineering network which is bringing together all the engineers into communities of practice. So traditionally they've spoken only within a project or potentially a

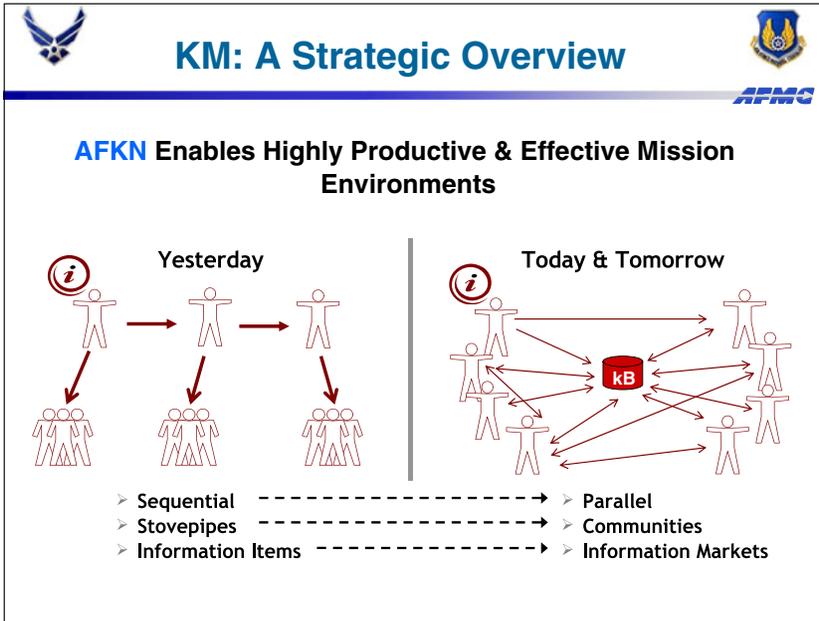


Figure 2. The Air Force Knowledge Now (AFKN) Strategy.

center...[I]t just doesn't happen [because] they're working on different types of projects and locations. But now this *new way* we're really trying to get these people to say "Look there are still new things we can learn between those projects from one engineering group to another." [emphasis added]

The KM Exchange itself (labeled as a CoP by many of our informants) is seen as bridging corporations in an extremely secretive and competitive industry hampered by nondisclosure agreements. One KM practitioner noted:

It seems evident that *all* of the groups that are participating in this community [KM Exchange] have pretty much the same kind of problem, and this is one of the *few* areas where we *share* the insights that we gain and the lessons we learn in the disparate communities. You know, like sharing stuff with [aerospace companies]...that is *very, very* unusual in today's environment.

There is a distinguishing discourse of doing "good" by forming communities of practice; as another KM practitioner mused, it is a moral imperative:

Bringing people together is a struggle...getting people willing to share. You know, especially in aerospace and in classified aerospace, especially, people have been taught not to share and I think people's nature these days is much more selfish and to get people to realize that by sharing, they're actually helping the enterprise.

This was an opinion echoed by many of our informants from the KM Exchange. CoPs promote a culture of sharing amongst engineers who usually work alone.

5.2. Roles in a CoP

For the KM practitioner, communities of practice contain *members*, some of whom have a specific *role* to fulfill. CoPs are part of a formalized process that requires sponsorship. More specifically, these formalized processes assign predefined “roles.” For example, one junior practitioner explained how her organization specifies myriad roles for each formal CoP:

Well, there’s certain requirements for communities of practice because we’re pretty strict...because we have different groupings. We have committees [and] communities of interest which essentially anybody can create, so those can be like the clubs or if I want to form a Unix group, then I could stick it in there, but communities of practice, they require a vice president as a *sponsor*, so someone has to say, “I’m accountable for this community, to achieve whatever goal that they’re going to achieve and so I sign off on those.” So they have to have their sponsor. They have to have a leadership team, so a group of people, like about 5 to 7 people, who will lead this community and, you know, keep it focused and be the ones who set the strategy. So they’ll have a *leader*, they’ll have someone who sets up the meeting, the *facilitator*, and they’ll have a *content manager* or slash *knowledge manager*, the one who’s responsible for filing their material, for keeping their space in our—well, we store all their stuff in Livelink. [emphasis added]

There are remarkable similarities across the roles embedded into formal communities of practice. In the Aerospace Corporation (Sutton et al. 2008), the defined roles are sponsor, leader, facilitator, content manager, information policy manager, topic coordinator, stakeholders, and members. For example, the facilitator “helps to ensure that the community dialogue and participation engage the full diversity of the member and stakeholder participants” (p. 97), whereas the topic coordinator is “responsible for leading the material development on specific topics overseen by the community” (p. 97). At the KM Exchange meetings we observed, numerous formal roles in CoPs were discussed. One corporation identified “content champions” for its CoP. Another had a management sponsor, community leader, and champions. For example, *people* champions “ensure quality and team spirit within the community,” whereas *knowledge* champions are responsible for “managing, building and enabling sharing of the community’s knowledge.” Another corporation appointed leaders, champions, and subject matter experts (SMEs) for each CoP. Finally, a government aerospace agency appointed a core team, human facilitator, and a leader. The leader “is a designated

Doing Business with Theory: Communities of Practice in KM

chief [Company Name] expert in the topic around which the community is centered” and is “tasked with considerable oversight responsibilities.”

On closer inspection, roles allow KM practitioners to ensure that CoPs work in a prescribed, effective manner to produce quality results. Whereas informal CoPs may not have monetary and official backing from the corporation, formal CoPs are beholden to their stakeholders to produce knowledge of the highest caliber. For example, such knowledge *produced* by a CoP is designated as “wisdom” by the Aerospace Corporation (Sutton et al. 2008). Concretely, wisdom might include best practices, design guidance, lessons learned, handbooks, guidelines, and feedback. This represents the retention of quality knowledge: “no CoP can be considered to have created knowledge *worthy* of publication as Wisdom [sic] unless input has been solicited from a broad spectrum of community members” [emphasis added] (p. 101). Thus, not only is knowledge objectified, but wisdom, the best of knowledge, also is objectified. When there is disagreement as to what is the official, representative wisdom of a CoP, the majority wins, and the minority views are retained but not made public. In contrast, communities of interest because they have “no constraints on the documentation of their knowledge or the processes by which they reach that knowledge,” have positions that cannot be representative of the company (p. 102).

5.3. Technological interdependence with CoP

Communities of practice are often spoken of in the same breath as collaborative technologies. Lexically, CoP and technology are used almost interchangeably in the following case study (“The Boeing Company, Rocketdyne Division”, 2003) of Boeing Rocketdyne (currently Pratt & Whitney Rocketdyne):

Boeing Rocketdyne Canoga Park has a large variety of communities of practice. They vary in complexity of issue attacked and methodology of communications used. Some communities just use shared drives; others have generated Web sites or have e-mail distribution lists. The more complex communities have turned to complete virtual collaboration tool suites for communications and are trying to generate KM as an end result of their efforts.

Interestingly, the spectrum of communities in terms of complexity is intimately *tied* with complexity of the computer-mediated communication utilized. Virtual technologies provide a means to create CoPs that join geographically separate groups.

Content management systems such as Livelink or SharePoint were commonly cited by our informants. As one practitioner described her role as content manager, “I...serve [in] a consulting capacity helping the members of the communities of practice use the Livelink tool as a document training tool to house their information and be able to retrieve it.” At the Aerospace Corporation, the “Knowledge Steward CoP” is in fact a community of practice devoted to setting up communities of

practice. This is a common activity found among all the aerospace companies we investigated—KM practitioners “eat their own dog food” to provide a poster child of how useful CoPs are. These KM practitioners ensure that knowledge, information, or wisdom is properly stored in such systems. Moreover, specific tools are an integral part of *implementing* a CoP, as shown by the following subset of a *required* 14-step process to establishing communities of practice at an aerospace defense contractor:

Implement and Build Community

- STEP 4 – Establish Livelink Workspace.** The Livelink workspace is the Community of practice knowledge Repository.
- STEP 5 – Establish community taxonomy.** A logical taxonomy is the cornerstone for building content that is useful and intuitive for the users....
- STEP 6 – Organize the content.** Once you have created taxonomy and determined where the knowledge resides externally, you must decide where each contribution belongs internally within the community....
- STEP 7 – Identify and develop any content engineered specifically to support the community...**
- STEP 8 – Identify Content Champions for the community.** Content Champions are responsible for monitoring both existing and new content: reviewing it, approving or denying it, featuring items to draw attention to them, checking for outdated or inaccurate materials, etc.
- STEP 9 – Train Content Champions on the KM tool.** To perform their job effectively, Content Champions must be trained how to use the KM tool. In addition, they must learn the basics of content management.

At NASA, one goal of its five-year KM strategy (Leonard and Kiron 2002) is to “support communities of practice through electronic and traditional processes.” This is accomplished by “developing Web-based collaborative environments such as a document manager that has an action-item tracker, a calendar, a team directory, a threaded discussion tool, and an activity log.”

This emphasis on web spaces illustrate how integral technological artifacts are to CoPs. To set up a CoP is to enroll not only procedural processes involving humans (e.g., establishing roles, planning community stakeholder workshops, and so forth), but also the mobilization of a web space that represents and encapsulates the knowledge created, produced, and shared as a *product* of the CoP. The web space reifies a taxonomy in which knowledge will reside. Thus, how the technology *should* be used in concert with the CoP is formalized.

Part of a KM practitioner’s survival depends on this unique set of competencies he or she potentially offers organizations. Not only do KM

Doing Business with Theory: Communities of Practice in KM

practitioners understand people, they understand how to use technology to implement particular types of communities (Su et al. 2007).

5.4. Evaluating communities of practice

As CoPs have become real objects of inquiry, they have also become accessible to rigorous evaluation. Some CoPs fail, while others succeed. Some are floundering and need help. Others are invigorated, having dynamic and active conversations and act as models for others to emulate. Those CoPs that die may be “resurrected.” KM practitioners recognize that hard, quantitative metrics are an imperfect measure of the impact of CoPs on an organization, but also know they must unequivocally demonstrate the effectiveness of their discipline’s tools to upper management (Su et al. 2007).

Metrics are derived to measure a “community’s progress” (Sutton et al. 2008, p. 103). A community’s “health” consists of measuring both the progress in building its infrastructure and the activity within a community. For the Aerospace Corporation, a CoP’s value is how far along a community has gone with respect to the standardized procedure of constructing a CoP (e.g., writing a charter, selecting a leadership team). Raytheon conducts a “performance and cost benefit measurement process” on CoPs. These assessments are conducted by a third party, and one of Raytheon’s presentation slides proclaims its CoP scored in the “top percentile” (Palla and Ward 2006). The American Productivity & Quality Center (APQC) plays an important role by providing a “health” assessment process for CoPs that focuses on “structure, leadership, knowledge sharing processes, communication, recognition, measurement, and outcomes.” The *Communities of Practice Implementation Guide* for one corporation states that a “CoP tool” tracks quantitative metrics such as “New Members,” “New Knowledge Objects,” and “Page Views.” Also outlined are measures of a CoP’s effectiveness through qualitative methods such as “conference surveys,” “workshops,” and “focus groups of users.” Many of these metrics are, in fact, *metrics of the technological artifact*, or web space, that is coupled with the CoP.

CoPs are also a *countable* set, and KM practitioners tally up communities of practice. How many does your organization have? One company counts those that are corporate-wide, cross-sector, and sector-specific (e.g., material sciences). One presentation by the same company noted that they had 68 CoPs. One KM practitioner told us they had 13 or 14 communities of practice, most of which are “technical” (i.e., formed around some subject matter).

Evaluations are also done in an iterative fashion throughout a CoP’s life cycle. This creates a continual *consciousness* of the CoP for both the members and the molders of a CoP. For example, the *Community of Practice Charter* for one particular corporation provides an *Early Progress Checklist* for CoPs. It asks questions such as whether the CoP has a common purpose, whether an appropriate leader is in place, and whether the proper technology is usable. All

these questions continually probe the CoP and make sure that it is progressing according to the KM practitioners' vision of what a CoP *should* be.

5.5. Discussion

Many of the informants related how difficult it was to disseminate communities of practice (Wilensky et al. 2009). However, KM practitioners believe CoPs are a humanistic process necessary to solve the disparate attributes of aerospace employees. They could potentially herald the creation of a new, novel environment of sharing.

In the organizational landscape, KM practitioners see formal and informal instantiations of CoPs lying aside other social forms of grouping, such as teams. Certainly, it is not our intention to imply that KM practitioners simply see CoPs *only* as objects to create and manage. They *do* recognize that there are trade-offs between the formal and informal, and that technology is not a panacea. One KM practitioner explained how she could direct people to CoPs for expertise (expert seeking is a common thread in KM):

In a couple of cases, I found existing CoPs that I was able to point these two guys toward and they were able to go to them for assistance. And [in] another case, there was no community or formal community. There were experts around the company...you know, just let them be [an] informal thing until they are ready to make them more formal.

Hence, CoPs are, necessarily, not (1) automatically transformed from informal to formal, and (2) created from scratch. At one KM Exchange meeting, a practitioner emphasized that they are always balancing structure vs. freedom.

Nonetheless, formalized CoPs remain as *the* site in which KM practitioners can best demonstrate their own aptitude and necessity. This opportunity allows KM practitioners in the aerospace industry to institute formalized processes, technological tools, and metrics in building a CoP. Although other CoPs may exist (e.g., one informant described a *taxonomy* of CoP: open communities, intra communities, and closed communities that require you to apply for membership), formalized CoPs are accorded special distinction. KM practitioners strive to create a corporate environment that recognizes CoP membership as a prestigious honor.

The issue of membership is further formalized via roles. We have seen that there is a class of explicitly defined roles for a CoP. This guarantees that formalized CoPs will produce "quality" knowledge, and it shifts responsibilities to particular members so that they will not abandon a community of practice.

As an object of inquiry, CoPs become receptive to evaluation, and once again KM practitioners are well positioned to provide expertise. Metrics, often derived from usage logs of technological artifacts, provide one means to see whether a

Doing Business with Theory: Communities of Practice in KM

CoP is “effective.” An iterative process, KM practitioners are constantly involved with keeping CoPs in good “health.”

In sum, the dominant discourse of a CoP in the aerospace industry is of a highly formalized process in which knowledge can be effectively captured and shared. It is framed as a *modern* KM tool (in contrast to such outmoded tools as databases) in that it takes people into account.

6. The Institute for Research on Learning: the locus of translation

Having seen how KM practitioners perceive communities of practice, we now ask: How is it that the concept of CoP came to become such a concrete, formalized theory? How did CoP enter the business world and become a de facto part of the knowledge management practitioner’s oeuvre? We argue that the Institute for Research on Learning (IRL) served as the locus of translation for CoPs into the KM discipline. In other words, IRL served as an intermediary for popularizing academic work.

A 1986 press conference with the CEO of Xerox hints at motivations for IRL’s conception (IRL opened its doors in 1987). It documents how IRL was formulated as a response to a *crisis*:

American business will have to hire more than a million new entry-level workers a year who can’t read, write, or count. Teaching them how, and absorbing the lost productivity while they’re learning, will cost industry 25 billion dollars a year for as long as it takes....[R]emedial training in the workplace is now part of the cost of doing business in this country. It’s a cost I resent, because when business has to teach basic skills, we’re doing the school’s product recall work for them. [David Kearns]

Crisis as a rhetorical strategy has been employed on numerous occasions; in human-computer interaction (Cooper and Bowers 1995), crisis elevated the “user” from a mere automaton. Moreover, in this promotion for IRL lies a causal relationship built on the idea that *education influences business*. Businesses are wasting countless billions of dollars in fixing something that our schools should have been doing. In response, Xerox heroically took charge of the matter by investing five million dollars to jumpstart IRL. IRL’s solution would become situated learning.

7. IRL as “spokesperson” for situated learning

Jean Lave’s tenure at IRL was early and short, but her presence had lasting effects. John Seely Brown recalled that “bringing Jean in [to IRL] was a discontinuity if not a radical disruption. And [we] had no idea how disruptive, you know, in a productive sense she was going to be.” The mantra of

“situationism” made IRL an institute dedicated to a social theory. IRL became a revolutionary institute that went against mainstream academia and business.

For Lave, it was meaningless to say that one can separate mind and body. If one were to talk about learning, one cannot ignore the social world in which learning happens. While such a belief may seem innocuous today, given the prominence of other research on socially situated activity (Lave 1988) (e.g., Suchman 1987), it is important to understand what the prevailing thought was during the late 1980s to the early 1990s:

Cognitive science, and computer science as well, is based on the absolutely pristine separation of mind and body, person and social world. The whole concept of situated learning violates those boundaries. It is intended to do that. You cannot separate people, as they move around in the world, from the world of which they are a part. This claim challenges assumptions underlying cognitive theory....I thought you should look at people engaged in the world if you wanted to understand how minds are part of bodies, part of their activities and part of the world. [Jean Lave]

By unpacking apprenticeship in the world, *Situated Learning* (Lave and Wenger 1991) provided a thoughtful blueprint for analyzing any educational practice without asserting that other forms of education were simply the opposite of institutionalized schooling.

Analogous to the notion of learning through institutionalized schools, corporations (and laypeople) largely equated learning with *training* (as evinced by the 1986 press conference in Section 6). Training was the responsibility of human resources (HR):

The Institute was remarkable in how much we did to get beyond kind of the HR and training world to really have direct operational line of business kind of impact....Early on a lot of our projects were in the training division....The learning world in the corporation can tend to become kind of...ghettoized and they have to fight for their budget dollars and they're trying to show their value more broadly....But they keep getting pushed aside as this nice little stuff that you get taken care of one week a year, and it's not that [simple]. [Melissa Cefkin]

Hence, IRL was able to effectively capitalize on situated learning's expanded scope. IRL sought to establish learning as a first-class object. Learning was important. Learning is not *limited* to a single, physical locale, but is something that permeates throughout an employee's everyday practices at work. By expanding the scope of learning to cover the entire corporation—not just HR—IRL positions itself as being even more vital.

IRL soon held the torch for situated learning. It became a force to spread that notion of learning. However, how did communities of practice (as opposed to legitimate peripheral participation or situated learning) become the focus of IRL?

Doing Business with Theory: Communities of Practice in KM

In the following two sections, we discuss two key steps toward the transformation of the CoP into a stand-alone theory.

8. Situating CoP for the corporation

In Duguid's (2008, p. 4) retrospective on communities of practice, he apologizes to Lave for an article (Brown and Duguid 1991) he wrote with Brown:

The link from Lave and Wenger's theory to Julian Orr's (1996) study of the Xerox 'reps,' going from...communities to the heart of corporate America, is made neither in Lave and Wenger's work nor in Orr's book, which studiously avoids a discussion of Lave and Wenger. Rather, it was made for better or for worse in an article I wrote with John Seely Brown, thus I have to take some responsibility for this redirection of the theory.

For many organizational and management science scholars, their first encounter with communities of practice is via Brown and Duguid's (1991) paper. Duguid (2008) notes that some authors have failed to understand that Lave and Wenger's monograph was a reaction against conventional theories of learning grounded in cognitive science. People have come to ignore the origins of CoP. Instead, they gain their knowledge by "proxy":

You can usually tell whether anybody writing organizational theory had actually read any...Lave and Wenger. Because if they refer to the Xerox copier technicians, you know they probably read Brown and Duguid, but haven't read Lave and Wenger. [Paul Duguid]

The two authors of the paper (Brown and Duguid 1991) have dramatically different takes on the motivations for the paper. For Duguid, the paper addressed how Orr lacked a theory of learning and how Lave and Wenger lacked a theory of organizations. Duguid (2008, pg. 5) thought "[b]ringing together the two—the learning theory...and the workers with their complicated relationship to the corporation as the dominant source of what they should know...seemed a worthwhile thing to do." Duguid seemed to view the paper as something innocuously conceived. Moreover, as a nonacademic at IRL (he was hired for his experience at a publishing house), he felt little need to popularize his paper as a traditional academic would:

So, I was just out of touch with where it [(Brown and Duguid 1991)] might be spreading...I didn't go to conferences. I didn't really move in the world. I didn't go to things like the Academy of Management. I didn't go to anthropology conferences. I didn't go to educational conferences.

When told of his immense influence by an academic from a prestigious business school, Duguid recalled that "not only did I not know what my influence was, but I didn't know what area she was talking about...and I was too embarrassed to ask

her.” After he had left IRL behind, the paper’s “relative celebrity” caught him by surprise.

In contrast, Brown viewed the paper as a key step toward constructing a *usable* theory. One cannot discount the effect that John Seely Brown had on popularizing communities of practice. A well-known speaker with a charismatic personality, Brown was perfectly positioned to create an influential institute. Inexorably tied with IRL was Xerox PARC. IRL was akin to a scientist hired to discover what makes Xerox PARC tick:

We were...actually trying to understand: What really goes on in innovation? What makes PARC actually work? How can I make PARC work better?...At PARC people would say...the ideas aren’t getting out...So, it became to me an incredible interesting sandbox to look at. [John Seely Brown]

And while Brown had experience as a professor, he noted that what formed the germ of the paper was not a result of a concerted effort to publish new research that would make reverberations in academia, but rather a *by-product* of answering the question of how Xerox PARC innovated. Brown remarked, “The irony was we didn’t set out to do any academic research.” Not only were Brown and Duguid trying to understand innovation, but, indirectly, we believe they were trying to legitimize Xerox PARC as a viable entity in and of itself.

Just as IRL was to solve what academia failed to do in education, Xerox PARC was to solve what academia failed to do in basic research. Xerox PARC promised real pragmatic solutions that were grounded in research. At the same time, Xerox PARC was neither a traditional university nor was it a traditional business. Instead, it sought to straddle both spheres. John Seely Brown takes pains to note the benefits of Xerox PARC over academia:

PARC itself was an amazing place that...had...more cross disciplinary work...because unlike a university, we really worked jointly on problems. And we didn’t have real silos...[I]f you actually notice in the kind of research [done] in universities, there is no craft anymore. I mean you don’t get tenure for craft...and so they’re [academics are] just individual thinkers in a way...You didn’t get promoted at PARC by writing papers. As a consequence, PARC produced a lot of papers that have had overwhelming impact...[Y]ou didn’t write a paper until you...thought you had something that was *really* good.

Brown’s trajectory toward his seminal paper with Duguid was to legitimize Xerox PARC as a locale where research could be applied in the real world.

Not only did IRL foster innovative research that was counter to what most academics and businesses thought of learning, it also helped to legitimize the activities of industry research labs. It would be a major breakthrough to find a general way to pragmatically capitalize on *knowledge* from an organization.

Doing Business with Theory: Communities of Practice in KM

Hence, what we see in the Brown and Duguid paper is a *transformation from an analytical lens into learning to a theory that is prescriptive*:

[W]e don't want to just construct an analytic category [of] which a community of practice is or LPP [legitimate peripheral participation] is. We really want to see *what work can it do...*[W]hat does it enable us to do that it didn't enable us to do before and in that sense what is the value of this? [John Seely Brown]

Situated learning and, more importantly, CoP provided a way to create a successful hybrid of academia and business. For Brown, it was “a whole new set of lenses to look at how to think about structuring a research center that actually creates brand new ideas but has a *fighting* chance of getting them appropriated outside.” The paper, he contends, wasn't simply addressed to academics but to the real world writ large.

Thus, one of the primary ways in which CoP entered the business discourse was through Brown and Duguid's (1991) influential paper. Lave acknowledged the paper's success, stating:

Conceptions of “situated learning” and “communities of practice” would not have become so popular if John Seely Brown and Paul Duguid had not published their paper [in Organization Science].

The paper served as a proxy for businesspeople to understand what CoP meant to *them*. While the paper's utility can be seen equally from both Brown's and Duguid's perspectives—that of Brown has taken hold. Digging deeper, we see that the paper was not a direct effort to influence academics or to establish a new descriptive theory of learning. Rather, the paper was meant to help legitimize Xerox PARC and therefore translate Xerox PARC's success in innovation (as revealed through the CoP of Xerox techs) into something generally applicable for any organization. Its influence on academics might almost seem an afterthought. This showed how “research-in-action” was a viable process for combining the best of academia and industry. This prescriptive maneuver allowed one to further extract CoP *from* situated learning and ask, “What work does the theory do?” Next, we will see that Brown and Duguid (1991) were not alone in carrying CoPs into the spotlight.

9. Owning a theory

An interesting question arises when we talk about situated learning or communities of practice: Are they theories? What does it mean to be a theory? For Lave, CoP (Lave and Wenger 1991) was meant to shape the *questions* one asks about learning:

The interesting thing about “communities of practice” was the idea that people engage in very heterogeneous ways in producing something and at the same time are able to re-produce their possibilities for continuing to do it. How does

that happen? That was the question...you know, every productive activity, every productive system has to do two things: it has to produce what it's going to produce and it has to produce its own continuation. The concept "communities of practice" was trying to point to the second, especially to ask "How does the continuation of a practice come about when the practice is not uniform or simple?"

Thus, CoP was to engender a scholarly inquiry into how groups of different people engage in activities that allow practice to be produced and reproduced over time. In our interview, Lave never refers to situated learning or communities of practice as theories; instead, she calls them "concepts" or "notions." Others at IRL have called it an analytic perspective.

Brown and Duguid's appropriation of Lave and Wenger's concept was pivotal in shaping and marketing CoP, but Wenger himself played a crucial role in solidifying CoP by inheriting it. By assuming a position of authority, he was able to elevate CoP during his tenure at IRL. By the time *Situated Learning* (Lave and Wenger 1991) was published, Jean Lave was no longer at IRL. One possible scenario was that situated learning would simply have become a well-known monograph passed around only in academic circles.

However, upon graduation, Wenger approached his advisor to ask for ownership of CoP:

When he [Etienne Wenger] finished his dissertation, he needed a focus for a consulting business and he asked me if he could "have" communities of practice...I said "of course!" He took it and ran with it. His years of hard work, along with Brown and Duguid's work surely help to explain why these ideas are widely available in the management community. [Jean Lave]

Susan Stucky explained when Lave left IRL, "Etienne sort of took over the mantle. And she [Lave] said [to] take care of Etienne, so, I did, but...he's also responsible for [it] going off [to be] sort of more prescriptive."

We see that the concept of CoP does indeed have some corporeal metaphor. Like a tangible object, it is passed onto Wenger from Lave. And, understandably, CoP becomes a way to make a living for Wenger:

The association of my name with the *brand* of communities of practice allowed me to make a living consulting on my own, you know, that was significant, there's no question. [Etienne Wenger, emphasis added]

Now sanctioned to expand on his dissertation work with Lave, Wenger (1998) went on to write his magnum opus, *Communities of Practice: Learning, Meaning, and Identity* (see Section 2 for a review) at IRL. Here, we see Wenger making a crucial jump in reifying CoP:

In the 1991 book we are focusing on what we called legitimate peripheral participation. It's a learning process....There is a community of practice and in that

Doing Business with Theory: Communities of Practice in KM

context, learning takes the *form* of legitimate peripheral participation. I think that intellectually, one of [the] important thing[s] about that 1998 book is that the communities of practice itself becomes the foreground, and it doesn't take... communities of practice for granted. In that book it says, "How do these form? How do they evolve?"...*once the community of practice [is treated] as a **container**, if you wish, for a history of learning is foregrounded then you can ask yourself the questions of, "Can those things be cultivated?"* [Etienne Wenger, emphasis added]

As a container, CoP became a distinct object of inquiry. Thus, through IRL's patronage, Wenger made CoP stand on its own.

With the *theory* of CoP now at the forefront, IRL arrived at its stable instantiation. Thus, CoP becomes not a tool for situated learning, but a tool itself. While not always explicit, CoP became a driving force in shaping IRL's mandate. We might imagine the theory of situated learning shifted to another concept-turned-theory (i.e., CoP). But, the latter actually became something very real and indeed almost "technology-like." The work of shaping CoP continued unabated even without Lave's participation.

"Communities of practice" became a vital component of IRL's discourse. By the time they (even early hires) arrived at IRL, many of our informants described CoP as taken for granted. It simply existed, no questions asked:

Honestly, people didn't come there [IRL] who weren't interested in the idea [of CoP]. [Shelley Goldman]

Well, at IRL, it [CoP] was just there....It was a part of [the] ongoing conversation. [Gitti Jordan]

I mean when I came there...I had a sense that it [CoP] was a part of [the] tool kit. [Charlotte Linde]

Not only was CoP wholeheartedly embraced at IRL, it was also being continually tweaked. Shelly Goldman noted that "there were a lot of conscious efforts going into forking on those ideas [of CoP], and *solidifying* them and making them understood by other people" [emphasis added]. Before Charlotte Linde arrived, she could see that "the institute was already thinking about it [CoP], working with it, pushing around to see what could be done with it." For example, Susan Stucky described how researchers "did posters and hands-on sessions for all our projects and all our efforts which we reinterpreted what the connection was through the lens of communities of practice."

Despite CoP's sometimes overbearing influence, Melissa Cefkin and others acknowledged the power of having a unifying theme, something that everyone could comprehend, complain about, and praise:

I've never worked anywhere that had quite such a unifying theme. I also recognize, organizationally, it really gave us a shared...something that we could have in

common, and just view or talk about or roll our eyes at....It was a shared reference point for all of us....It became this external third thing out there that between any two of us we could at least know we had that in common. [Melissa Cefkin]

Having passed from Lave to Wenger, CoP provided a set of useful discourse for employees at IRL to communicate amongst themselves and others.

10. Marketing the new CoP mindset at IRL

CoP thus became the flagship product of IRL:

People *came* to talk to you because of those ideas [CoP]. [Shelley Goldman]

That's what we were known for [CoP]. That's why clients came to us and that's how the conversations initially with clients often started. [Melissa Cefkin]

Concerted efforts in marketing itself as an innovative organization that mixed both academia and business allowed IRL to position itself as offering a true theory with application. We also see how IRL reified CoP to make its unique services more palpable to potential clients.

Part of IRL's success with CoP was its innovative structure. Its nonprofit status gave it a pristine, unbiased image. A section in an early 1988 pamphlet entitled *The Advancement of Learning* attempts to clarify common misconceptions about IRL. In this blurb, the pamphlet makes it explicit that IRL is not your usual corporate machine. IRL will not bring "corporate products into classrooms." It is not a "commercial" manufacturer. It is an independent and nonprofit institute whose results will be disseminated publicly.

Thus, for Xerox, IRL was a way to do a public good:

Well [what] they [Xerox] were standing to gain was publicity. And [it] seemed to be doing something good, that's what those foundations do. [Susan Stucky]

The nonprofit status of IRL attracted researchers, such as Cefkin, who still wanted to keep one foot in academia:

The fact that it was not for profit...mattered enormously to me...that sense of doing something good in the world...not being you know, motivated by a profit...or by a sense of sort of the competition of the corporate world or business world that sort of framed around that. That was just never my orientation. [Melissa Cefkin]

At the same time, IRL would accept only clients that allowed them continue to produce the primary products of academics, publications. Charlotte Linde remembered that IRL would not agree to "do corporate projects if we could not afterward do open publication."

Doing Business with Theory: Communities of Practice in KM

Even though IRL was a nonprofit institute, it still needed to make money to stay alive. Xerox's plan was to initially fund IRL with the hope that it would eventually be able to sustain itself from clients (this would prove unsuccessful). Thus, like any other business, it had to market its product, community of practice. Susan Stucky recalls first hearing the term "communities of practice":

Within a hallway on Hanover Street. It has to be at the end of the second year [1988]. And Jean Lave came down to the hall and said, "I have a name for it." And I said, "What is it?" And she said, "Communities of practice." And I thought, "Oh, God! How the hell are we going to sell that?"

While academics certainly promote themselves via presentations, publications, interviews, and the like, the style of promotion is a different beast in industry settings. Susan Stucky realized the importance of translating results for IRL to earn money:

This isn't fun and games...nobody thought about [the] logo...I was employee number three....and it was...I got there and I thought, "Oh my god! Nobody thought about stationery." And then there is an institution to build here.

As a locus of translation, marketing materials helped translate academic work to something business and management could understand. In its 1993 annual report, IRL is listed as offering to "help our partners 'see' the learning that naturally occurs in the everyday activity of the classroom and the workplace." This "seeing" is accomplished through methods, tools, and theories. IRL offered novel techniques from anthropology such as interviews, participant observation, and interaction analysis. Software tools aid the analysis of micro-level behavior. And, finally, the IRL report explains theory:

[T]heory functions as a set of eyeglasses through which we view each new situation....Our work in the classroom and in the workplace has required that we make a conceptual shift from seeing knowledge as a property of the individual—as a kind of quantity that can be measured, assessed and transferred—to seeing knowledge and meaning as socially constructed within what we call "communities of practice."

Indeed, this seems like a reasonable distillation of what role theory played in IRL. In this and other brochures, IRL consistently strives to emphasize that knowledge is not simply something transferred in and out of the mind like some blob in a database. To depict an antiquated theory of knowledge transfer, illustrations such as those in Figure 3 are prominently placed in much of the IRL promotional material. Crucially, such figures helped to translate the one part of what Jean Lave and Etienne Wenger were striving for in their *Situated Learning* work—that learning is best understood in the context of its production and reproduction. This is somewhat ironic in that knowledge management, which has so readily adopted CoP, nonetheless still retains much of the older view on knowledge (knowledge transfer, capture, and so forth are commonly heard)

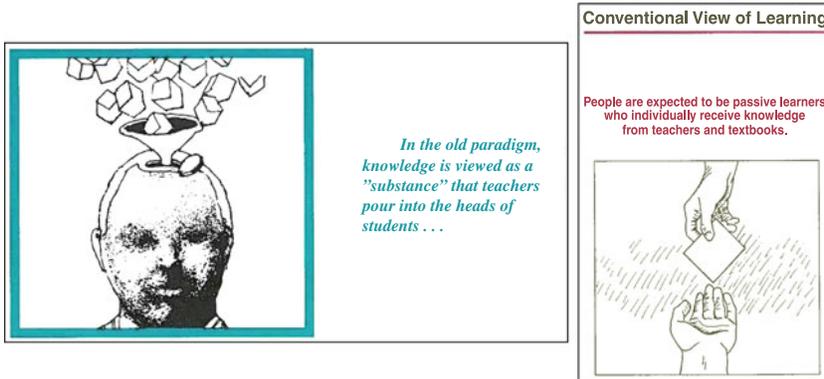


Figure 3. The antiquated view on knowledge transfer in learning.

(Su et al. 2007). Here we see that marketing materials emphasized that this new theory of CoP would replace outmoded theories.

Prominent among the marketing materials are the *seven principles of learning*. These seven principles were important in distilling CoP to business:

It [CoP] was the most *marketed* notion that we [at IRL] revolved around...At one point there was this seven principles that IRL developed that [had] communities of practice at the core, but the seven principles unpacked it for many businesses or lay audiences. [Melissa Cefkin]

Although PowerPoint slides are often admonished for *distorting* science (Tuft 2003)—a viewpoint in concert with the dominant view of popularization—their ubiquity in corporate America (and in other domains, such as the military) points to their key role in translation:

I don't think our clients ever would have read anything like that [*Situated Learning*]. We were *translating the notion* for them...the assumption was that probably everybody at IRL had read it...but beyond that I don't [think] we assumed people were reading it. And, in fact...in terms of the way these things move, it's not that Jean Lave or even Etienne's name necessarily remained attached to it. [Melissa Cefkin, emphasis added]

The seven principles of IRL (first found in a 1990 brochure) are:

1. Learning is fundamentally social.
2. Knowledge is integrated in the life of communities.
3. Learning is an act of participation.
4. Knowing depends on engagement with practice.
5. Engagement is inseparable from empowerment.
6. "Failure" to learn is the result of exclusion from participation.
7. People are natural lifelong learners.

Doing Business with Theory: Communities of Practice in KM

In our archives, detailed explanations are printed under each of the seven principles. These explanations draw extensively from communities (of practice). For example, principle 5 notes that “individuals perceive their identities in terms of their ability to contribute...to a community.” Principle 7 claims that people have a need for “meaningful participation” in communities of practice that motivates people to learn.

Another version by Gitti Jordan in the 1996 *Xerox Business Service (XBS) magazine* featured *eight* principles. This later version of the principles has a distinctively more prescriptive turn to it. For example, “Learning *needs* a supportive environment [emphasis added].” In addition, it is more specific, stating that “Learning crosses hierarchical boundaries.”

In practice, what did this new vantage point engender? What was IRL’s product? How does one “sell” CoP, a theory? One of the key services provided by IRL was “making the invisible visible”:

It [CoP] makes a lot of sense...people intuitively know...that the org. chart doesn’t reflect that [a CoP exists]. OK, so, it was very anthropological. We were arguing what we can do is...make the invisible visible by analyzing a community of practice other than going down [an] org. chart and asking survey questions. [Gitti Jordan]

A pamphlet from 1990 under the subheading of “How We Work,” touts:

We locate the unofficial and informal “communities of practice”—the true carriers of learning—that exist within the official groupings of the company.

CoP became an object where learning *actually* happens. It is where learning is *carried* (note how such words serve to give learning itself a concrete form). Examples shared by our informants about making the “invisible visible” include discovering how insurance agents are motivated not by extrinsic rewards but by the intrinsic reward of being part of a CoP with other insurance agents in a company. Another example includes discovering how a “super” performer in accounts receivable (labeled the famous “Carlos” knowledge management case study) drastically increased the productivity of those in the periphery of the CoP of call center agents.

The *identification* of communities of practice by IRL as a service also had the advantage of allowing one to compare and contrast *multiple* communities in an organization:

So, when clients came, people came out from woodwork and they did. And there were little communities of practice....There were people who were brokers who were kinda connected...[They] had some participation in more than one communities [sic]. And Etienne defined [a] configuration on this. This constellation of communities of practice. *Some were more identifiable than others. Some were more robust than other[s]*. [Susan Stucky, emphasis added]

Not only did IRL unearth the CoP specimen in an organization, it also could claim unique expertise in being able to characterize the *type* of CoP to which a specific, discovered CoP belonged. Both steps served to further reify CoP.

Translating CoP through brochures, posters, pamphlets, slides, presentations, meetings, and so forth served not only to entice clients to hire IRL, but *in turn also to instill in clients a new set of discourse*. Etienne Wenger would give workshops to clients to explain the merits of CoP. It was a discourse that was easily digestible for those in business and, as a result, they adopted the new language:

It was so much of how we marketed ourselves and spoke about the value in our approach, [that] our clients also very much...started to use the notion of communities of practice. So even they were talking about it and referencing it and stuff. [Melissa Cefkin]

Cefkin recalls how shocked her former colleague, an anthropologist, at IRL was to hear CoPs being uttered by everyday business people:

[T]he first time we were like in a business meeting with clients...and she heard them talking about communities of practice, she was *completely* floored. She could not believe that this language had entered the corporate world itself and they were using it.

The Institute for Research on Learning served as a locus for translating CoP—an “obligatory passing point” (Callon 1996). It mobilized business clients by posing itself as a solution to a learning/training crisis that had been inadequately addressed by academics, industry, and government. As a hybrid, nonprofit institute out to do a public good, IRL could be a trusted actor. Press releases and advertising material strove to portray IRL as a revolutionary, innovative institute dedicated to furthering learning—a goal that is difficult to criticize. Such materials announced IRL’s services surrounding CoP that further reified CoP into something that could be counted and classified (something Lave did not anticipate). However, we do not mean to create the illusion that this translation was unproblematic. In the next section, we see how the knowledge management discipline translated IRL’s own translation.

11. Translating the translation: KM’s adoption of CoP

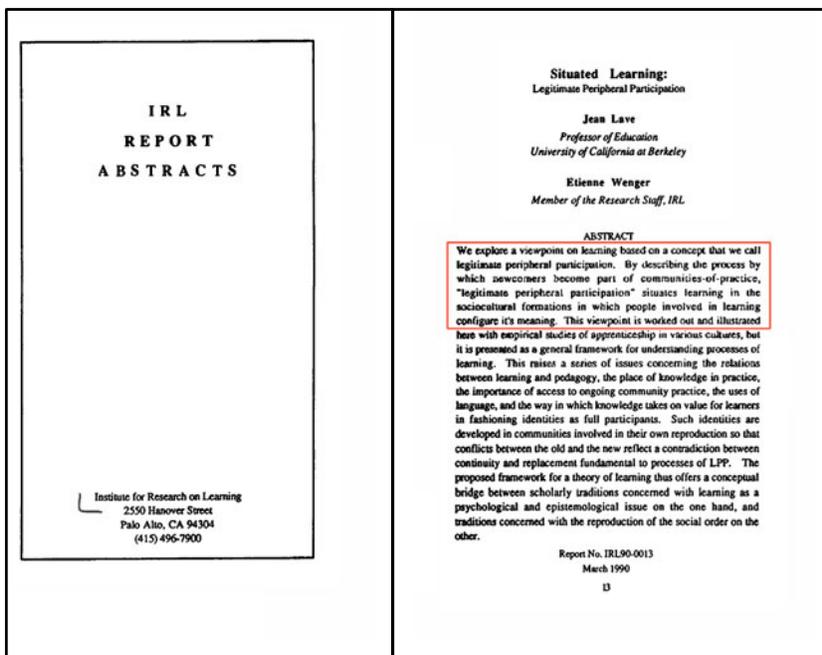
IRL was hugely successful in circulating the mantra of communities of practice. Rare as it is to have an entire institute dedicated to popularizing an academic theory, concept, or analytic lens, what made CoP so attractive for knowledge management practitioners? In this section, we discuss the lexical

Doing Business with Theory: Communities of Practice in KM

properties of CoP, its interpretive flexibility, its lack of “policing,” and its fortuitous timing.

11.1. Loving “community”: the lexical properties of CoP

What’s in a name? The *lexical* properties of a theory are important and contribute to much of its successful popularization. Yet, how we choose to name our ideas, theories, and inventions can have radical consequences on their popularity. As it stands, “communities of practice” is composed of three commonplace words: “communities,” “of,” and “practice.” Interestingly, however, communities of practice as referred to in the first tech report issued by IRL is written as “communities-of-practice” (see Figure 4), with hyphens.



ABSTRACT

We explore a viewpoint on learning based on a concept that we call legitimate peripheral participation. By describing the process by which newcomers become part of communities-of-practice, "legitimate peripheral participation" situates learning in the sociocultural formations in which people involved in learning configure it's meaning. This viewpoint is worked out and illustrated

Figure 4. February 1990 IRL Report Abstracts: Situated Learning abstract.

Duguid (originally hired as the editor for IRL's works) recalls that the hyphenation was not just for aesthetic reasons:

There was a discussion early on...I think possibly [in] the first IRL report...the term is hyphenated. And the idea was that it couldn't be broken up. It wasn't just about communities or just about practice...but I don't think it [the hyphenated form] made it into the book [(Lave and Wenger 1991)]. But that the concept was unitary and there was sort of some concern that it would be broken up [and] people took hold of *community* and for the most part left *practice* behind.

We do not know for sure what decision process settled on the unhyphenated form. Nonetheless, the fear that it might be separated was well-founded. In time, what we find is that the community component becomes prominent, and the practice component goes by the wayside.

As opposed to legitimate peripheral participation, the phrase communities of practice has an intuitive feel to it. Referring to Williams (1985), Duguid (2008) describes "community" as a warmly persuasive word. John Seely Brown explains the allure of community:

First of all, the notion of practice is a bummer. The notion of community is a plus. And, Paul [Duguid] is famous for saying at some point, "If we had called this a 'cadre of practice,' it would've sunk like a lead balloon." Although, a cadre of practice is probably a closer term...everybody glogged onto the community notion...we thought community was always a positive term. And so that was round, fuzzy, and good. You felt good about that.

Wenger acknowledges the baggage the "community" term has:

It is true that of course any term will have some baggage with it. So, the term community has baggages of peace, you know. Or of like loving, you know, or no power struggles....You know, good communities are full of conflict....They are full of people debating things and disagreeing....[A] nother problem with the term community is that...for some people it connotes homogeneity.

Amongst clients, Melissa Cefkin noted that the word community had a depth *beyond* simple new "process redesigns" that other consultancies would hawk, but still felt tangible and actionable. While communities proved attractive to upper management, it also threatened to bring about a shallow interpretation of the concepts IRL stood for:

You can go to the CEO and be able to say I'm doing this and he's going [to] say, "That's great because communities must be a great thing." But, you know, that...obscures a lot of what...you really need to be really paying attention [to] and what the value of [is] what it can do. [Melissa Cefkin]

Doing Business with Theory: Communities of Practice in KM

On the other hand, practice itself was a more difficult concept to grasp:

I don't know quite what [practice] is but it doesn't matter because the term community is so *cool* and practice has something to do with profession...[T]he focus was on the first term...and there was so much of a fascination and belief in the power of community that everybody overlooked the epistemological and ontological stances for what practice meant. [John Seely Brown]

We believe it unlikely that Lave even conceived that people would ever treat community and practice as separate concepts to be parsed and analyzed. Because of the dominating appeal of community, some would equate community with community of practice. This dramatically increased the *interpretive flexibility* of CoP. Its domain of application was expanded.

Community is certainly easier to depict pictorially as well. Much of the promotional material from IRL used graphics that depict people in a community, working together. And while the intent was not necessarily to downplay the inherent tensions that are always extant in a CoP as conceived by IRL (and Lave and Wenger), the images do paint a picture of happy cooperation, even featuring rainbows (Figure 5). These would be the very images interpreted by IRL's clients.

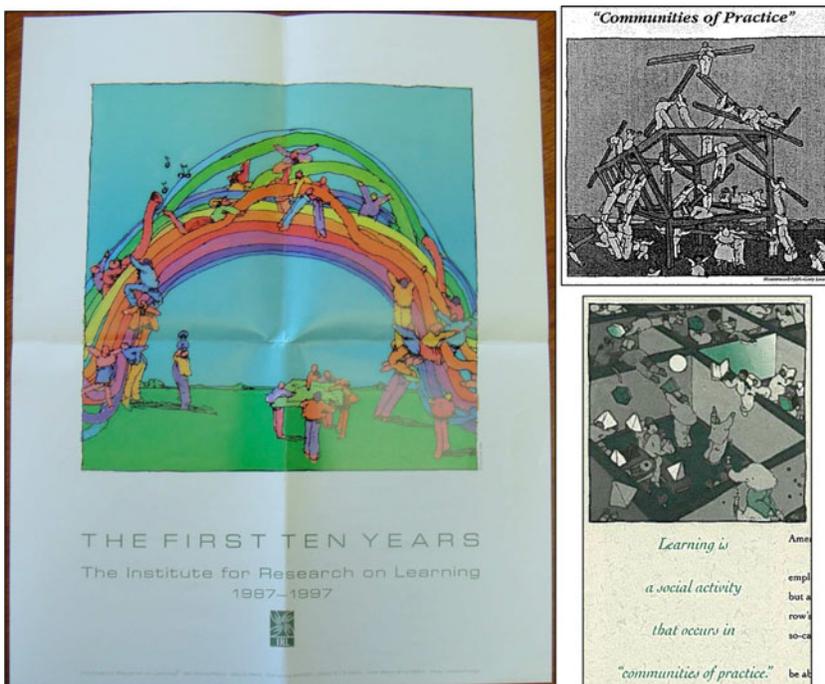


Figure 5. Various illustrations from IRL depicting communities of practice.

The lexical properties of “communities” in communities of practice partially contributed to its popularization. While the harmonious connotations of the word “community” has been mentioned in past research (e.g., Cox 2005), linguistic strategies such as hyphenation were utilized initially to keep communities and practice “together.” By interlocking these parts, it might have limited the scope of the term. Instead what has happened (certainly not just due to the lack of the hyphenation) was a separation of community from practice. As the more congenial brother of practice, community may have contributed to IRL’s success. For Lave, who remained firmly in academia, choosing the word “community” may not have been the best choice:

Using the word “community” was a bad idea, I guess. It certainly got taken up with the wrong understandings attached to it.

These wrong understandings as viewed by Lave and other academics are discussed in the subsequent sections.

11.2. Formalization of CoP

The analytic perspective of situated learning in which CoP is grounded argues that much of our learning is done in *informal* settings. Turning the informal into formal could very well be seen as a step backward into traditional models of learning.

Whether purposeful or not, the clients themselves are important in translating the translation of IRL. IRL translated the works of Lave and Wenger. The clients then translated what IRL told them. Like any actor in a network, these clients are indubitably affected by their own networks, past and present. So, those in upper management, trained in the “design” of organizations, will be influenced by their educational system’s frame of thought. Chief within this frame is an *instrumental* oriented focus:

I think business schools, like most professional schools have a highly instrumental side that [in] teaching MBAs, they’re often serving as consultants themselves....And in those roles, they’re called on, for better or for worse, to deliver solutions....But I think, it puts a *huge* amount of pressure on them to be instrumental and I think one of the things that probably distressed me most about [how] the communities [sic] of practice works is how it was instrumentalized very quickly. It...turned into a means for an end. You form communities of practice and solve corporate problems. [Paul Duguid]

The sort of circles the higher-ups at IRL traveled amongst were people who were trained in precisely the instrumental approach promoted by business schools, as described by Duguid.

Likewise, informants reported how clients immediately transformed CoP into a prescriptive entity. Susan Stucky quipped, “Oh, I think they [KM people]...get the idea it’s not formal...But then they immediately formalized it.” Informants

Doing Business with Theory: Communities of Practice in KM

explained their frustration that clients were focused on a very programmatic, objectified perspective on CoP:

Clients...were like, “Oh yeah, OK, well let’s figure out which communities of practice we’re going to organize, or have, or whatever,” wanting to start to list them right up front as if...it became a *goal* of something you’re trying to make rather than understand what was already inherent in the practice...but it was like, “Oh, we’re not doing anything over here, maybe we need a community of practice on that.” [The expectation was that] we’re going to set up these communities of practice. [Melissa Cefkin]

I remember how...especially in education...people thought you created situated learning. You had recipes, right?...It was almost like a recipe: if you did it you’ll have situated learning situations and people at IRL were thinking at a different level. [Shelley Goldman]

Such a perspective illustrates how organizations began to view CoP as a countable object. More specifically, organizations *without* CoP would lack the competitive edge, and such a “gap” needed to be “filled” by creating CoP objects.

Time and time again, our informants repeated with conviction that IRL was not out to manufacture communities of practice. While some interpretive flexibility was allowed for CoPs, this was where many of our informants had a clear demarcation on what was or not appropriate for the theory of CoPs:

Communities of practice, originally, was not formally, constituted....They used to say, you can’t make one, but, you can go out, do the research and find one or find several and then what corporations can do is to support them. [Gitti Jordan]

They [our clients] said, “Oh, we want to make, build communities of practice.” [If] this is how knowledge is embodied in organizations, then “Let’s have some.”...[I]t made me itch...I thought, “You got them. Let’s go and find them.” [Susan Stucky]

People starting saying, well, you know, hire me and I’ll come design you a community of practice. That’s how this term...really had lost all its analytic precision. [John Seely Brown]

Instead, IRL advocated supporting, nurturing, or facilitating communities of practice. The dominate phrase that Wenger carefully chose was “cultivating” communities of practice—the title of one of his books (Wenger et al. 2002). As we discussed in Section 10, one of IRL’s key competencies was in making the invisible visible. After uncovering the CoPs lurking in their organizations, IRL provided longitudinal services to capitalize on these “discoveries”: “Then, they [clients] had a choice. They could say, ‘OK...do we want to enhance this practice? Do we want to support it?’” [Susan Stucky]

John Seely Brown considers this “cultivation” a reasonable characterization of IRL’s work, but is aware that it is but a small step to translate cultivating into building:

Cultivating is a very carefully chosen word. *That got transformed into building.* He [Etienne Wenger], I think would be the first to agree...that you don’t build communities of practice. Communities of practice exist. What you may do is discover them, afford them more, and so on and so forth. So cultivating is an honest—I mean, it’s a fair word to use. [John Seely Brown, emphasis added]

While the employees at IRL viewed themselves (or wished to) as *researchers*, clients viewed IRL employees as *consultants*. As such, they expected that consultants would point out what was wrong, and furthermore that the consultants would be instrumental—that is, instigate change, not simply cultivate:

The corporation who hired us thought of us as consultants. And, one of the things that I haven’t realized, as a consequence of that, you hired consultants to give you bad news. If everything is going fine, you wouldn’t bother. You know, “This is business you got wrong. This is what you need to improve.” We couldn’t just give them happy news. Or we wouldn’t be [any] use for the company itself. [Charlotte Linde]

Indeed, organizations are constantly striving to innovate themselves. And sometimes change is done for change’s sake (Feldman and March 1981) to symbolize organizational progress. That is, CoPs promise a direct route to change. Charlotte Linde explained that many corporations had a “desire to get change free” through CoPs.

Despite such misgivings, IRL needed clients. They needed funding to support their employees. So, as a foot in the door, not totally denying the programmatic approach may have been appropriate if it allowed IRL to educate the business world on situated learning and CoPs:

I’m a little uncomfortable with it [formalization of CoPs] to tell you the truth. I don’t think we understand these communities enough to be so programmatic. But as an *entry point*, you know, for some people they need that kind of programmatic approach because otherwise they feel lost. [Etienne Wenger, emphasis added]

So while academics may scoff at formalizing CoP, the *pragmatic* turn to allow CoP to spread and to ensure IRL’s survival may have necessitated it. Yet, what does one sacrifice by entering a more prescriptive stance on CoP?

11.3. Losing analytic precision

Wenger is very mindful of the dangers of turning descriptive to prescriptive:

Yeah, do you lose your analytical sharpness when the tools that you use for analysis becomes a design intention?...Do you lose your ability to see these

Doing Business with Theory: Communities of Practice in KM

important differences between the formal organization and the informal organization when *the informal becomes the design of the formal*.

Indeed, for Wenger, this is not an issue of descriptive or prescriptive but descriptive *and* prescriptive. How can you maintain a hybrid of both in a theory?

It's *both* a design perspective and an analytical perspective, the fact that you say, "Oh, I'm designing a community of practice" does not relieve you of the need to be analytical about it and see, is this really what's happening here?" [Etienne Wenger, emphasis added]

The problem is that once someone starts to delve in either side of these two facets of theory, the other side may be ignored.

For some, CoP has already lost what precision it had through its multiple translations:

It lost its analytic precision. But...it may have become easier to apply. So, I think...he [Wenger] had to fight a lot of kind of old thinking in the organizational world and so...*when you're fighting a war you cut corners*. So, I mean, you know, to me, I suppose some of the theoreticians, you know, could pull out their hair over how they loosened the meaning of the term. [John Seely Brown, emphasis added]

One way to "save" a theory from being diluted is to propose a separate term that catches the diluted form of a theory before it destroys the original, precise form. Certainly, the number of terms invented to denote an alternative, narrower definition of CoP is perhaps an indication of how widespread was this belief that the equivalence of CoPs and communities was encroaching the space. By far, the most popular extension of CoP has to do with multiple communities. For example, there are communities of practice and knowledge (Klein and Hirschheim 2008), communities of interest (Fischer 2001), constellations of practice (Wenger 1998), and networks of practice (Brown and Duguid 2001).

Another thing that concerned me was the way that concepts of community of practice were used to talk about professions or academic disciplines. So people would talk about the community of practice of physicists, and that made no sense to me. And the way that I tried to resolve it was to suggest that there are *networks of practice*, which [could be] fairly significantly displaced geographically. [Paul Duguid]

I think the way that communities of practice has most often been applied is really what I would call more *communities of interest* where people affiliate based on something they're explicitly thinking about or know about. [Melissa Cefkin]

But community of practice proved to be such a successful term that instead of successfully migrating different notions of CoP to new terms, CoP assimilated any competitors:

As soon as you get networks of practice then you start getting now the more popular notion of a community of practice. Well, few people have accepted the framework of networks of practice. What they've really done instead is overgeneralized to community of practice. To be kind of these wide distributed communities that share a tiny bit of practice...so it's lost a fair amount of analytic precision. [John Seely Brown]

Wenger is careful to state that “when we talk about communities of practice, we are talking about a very, very specific kind of community...not just any old community.” LPP is similarly “a very, very special case of learning within a community which is when a new member is becoming a full member.” Just because a “group calls itself a community of practice does not mean that it is one.” Wenger here is striving to prevent the adulteration of CoP. Yet one must acknowledge that by making CoP broader and applicable to a variety of situations, it became an easier sell to clients.

Was a drift in the interpretation of CoP inevitable? One might consider that this loss of analytic precision speaks to the theory's malleability. Is it possible to even police a theory? Do the inventors have some sort of priority to do this?

11.4. Policing theory

Recently, Paul Duguid and Jean Lave wrote a prologue and epilogue, respectively, (Duguid 2008; Lave 2008) that provide retrospectives on the concept of communities of practice. Their writings reflect upon their own roles in the concept. Both Duguid and Lave acknowledge that CoP is an “active construct... susceptible to constant change and redirection” (Duguid 2008, p. 1). They make no attempt to publicly make clear a “correct” definition of CoP. While congenial in tone, both chapters ask that readers look at the origins of communities of practice. As Duguid said in his interview:

And I was trying to say in that introduction, it seems to me on the one hand I can't object to that range or say they're wrong but, I can sort of say where I thought I was going with the concept when I talked about it.

Lave (2008, p. 295) wistfully noted at the end her epilogue,

Clearly I believe our ‘take’ on communities of practice was part of a powerful and productive approach to social analysis. But it is also clear that this is now just one conception of communities of practice among others.

Doing Business with Theory: Communities of Practice in KM

To a certain degree, the originators and key players of CoP all have a conciliatory tone regarding the direction the concept, theory, lens, tool, etc. have gone. They cannot, nor will they, police the theory:

You know, I'm not going to write the person and say, "You're wrong. You can't say that."...When I realized how much my own understanding of it is evolving, you know, who am I to point my finger, you know? [Etienne Wenger]

If people want to use community of practice to mean a corporation, I can't stand up and say, "You're wrong" you know...Language takes its own life. [Paul Duguid]

Indeed, even for the inventors, theory itself is an active construct and can only be construed at points in time when it is stabilized in textual form. And it seems rather futile to control a concept once it reaches a particular point: "it got too broad to try to control and the more broad it got, the less interest or purpose there was in controlling it" [Paul Duguid].

We can only surmise why the theory is not actively policed, but it could be said that the grounding in social theory that many of our informants had gave them a perspective that any such construct that is part of the everyday is socially constructed (Berger and Luckmann 1967). Lave explained why social theory is vital for those who seek to understand and instrument the world:

I think the misunderstanding [of CoP] comes from the fact that cognitive and computer sciences are basically common sense positivist theoretical endeavors. Those fields do not include adequate education in social theory. Until more computer scientists and cognitive scientists get a serious education in social theory they will go right on distorting and assimilating social practice theory to positivist assumptions rather than working to understand social being on the basis of social *historical* relational assumptions. [Emphasis added]

Ironically, such a reflective perspective may have made the translation of CoP a self-fulfilling prophecy.

So translation occurred unabated without any sort of "central" authority or academy. Institutions such as IRL were able to incorporate and translate CoP for the masses. In some cases, misconceptions proved to be pragmatic steps to ensure that the CoP became a vital actor of the business and knowledge management disciplines. While we do not claim that such misconceptions were *actively* encouraged, their benefit to actors in KM is unquestioned.

11.5. Updating knowledge management

Direct interaction with knowledge management practitioners did occur via IRL. As the "birthplace" of CoP, IRL/Xerox PARC could claim some authority over the concept. John Seely Brown acknowledged this, "I think the knowledge

management movement based on communities of practice absolutely came out of PARC.” Susan Stucky comments on being part of the KM “circuit”:

I was on the knowledge management talk circuit for a while...Conferences... thought leaders, you know, [with] people like Larry Prusak and Tom Davenport...[T]here were conferences, the occasional keynote...John Seely Brown was doing a lot of that, as well.

Stucky also mentioned that KM practitioners from companies would come to IRL due to its expertise in CoP:

He [a KM practitioner] was the one who came and said, “We are in knowledge management. We[’d] like to introduce this concept, communities of practice and can you help me, you know, figure out the way to build [it].”

IRL’s other focus on education sometimes brought fortuitous connections with business because many people from corporations had children of their own and had interest in schools. Shelley Goldman noted that she met Estee Solomon Gray (who created positions such as Chief Knowledge Officer) through the board of a day school and eventually had Gray consulting with IRL.

Inevitably, these relationships with companies invested in knowledge management led to further opportunities. Sometimes it led to further business opportunities. For example, Wenger told us of his time at The Society of Organizational Learning conference organized by the Senge Group. He was approached by the Shell Oil Company:

That was just a fluke...one of the leaders of their knowledge management practice came to a conference where I gave a talk...And after my talk she said, “We need you.” So, I said, “Great!”

McKinsey & Company, a prominent consulting firm, had a number of prominent knowledge management professionals who interacted with IRL:

I became friends with some people from McKinsey who were well known in the knowledge management field, Nathaniel Foote and David-Brook Manville, and..this is [when] we started to talk about our work in knowledge management...and then...[CoPs] became adopted as part of knowledge management. [Etienne Wenger]

Tellingly, when asked to recall the moment he first realized communities of practice were catching on in the world, Wenger replied that it was at a knowledge management conference:

[T]here was a...knowledge management conference in Boston where I gave a talk and the reaction was like *so* positive that I said, “Wow, this is really *significant* to people.”...I think [the conference was] in June 1998 in Boston.

Doing Business with Theory: Communities of Practice in KM

What made community of practice so attractive to KM practitioners? It came at the right place at the right time. Like a knight in shining armor, CoP promised to rescue knowledge management from its creeping obsolescence. It modernized a discipline, giving it new life, without destroying those who had invested in establishing that discipline as indispensable to business (what Zizek (2008, p. ix) describes as “Ptolemization”):

The feel of knowledge management at the time was kind of like a bit stuck in technological approaches...but knowledge is about people and we don't talk about people, we talk about machines...[Y]ou would go to a talk and they [the speakers] would say, “Well, you know, knowledge management is 10% the technology 90% people.” And then the rest of the talk would be about technology because they didn't have very good ways of talking about people as carriers of the knowledge of the organization. [Etienne Wenger, emphasis added]

Those at IRL knew that CoP provided a paradigm shift for KM practitioners:

So, to suddenly have this concept of a community of practice allowed you to start thinking about knowledge management in a new way, you know. That knowledge management really has to do with recognizing and enabling these communities to exist, which is different from just creating knowledge flows and databases, you know. [Etienne Wenger]

And, knowledge management was coming in and saying we capture these nuggets of knowledge which never got used...and they're getting stripped out of the context. And, so we became much more sensitive to the role of context. And that was the parameter being misunderstood. [John Seely Brown]

Hence, the job of KM practitioners shifted from database-centric (capturing, storing, and transferring knowledge) to cultivating communities of practice.

Orr's (1996) ethnography of Xerox technicians particularly resonated with those in KM who were seeking to upgrade their discipline to account for people while still providing an instrumental approach to deal with that thorny people “problem”:

[T]he work Xerox then did on the Eureka project which involved providing a database of tips that were volunteered by the technicians...I think that appealed very much to knowledge management. It seemed, on the one hand, to look suitably responsive to the complexities of knowledge, rather than have an instrumental view of that...[I]t seemed responsive to the soft side of knowledge but simultaneously it offered a technological solution—a relatively low cost one and that by Xerox's account saved the corporation a great deal of money. [Paul Duguid]

Thus, CoPs put forth what Wilson (2002) terms the “people track” of KM. The war stories the Xerox technicians exchanged were transformed by KM practitioners into a more generalized technique of “storytelling” to spread knowledge in a community.

Eureka and subsequent technologies that are grounded in CoP are a result of a distinct coupling of technology with theory. For some, creating a community of practice can become creating a web space. What is interesting here is that the scope of CoPs has expanded significantly. Before, in Brown and Duguid's interpretation, CoPs were relegated to tight-knit, face-to-face communities:

[The] tightly focused joint work got done in these face-to-face communities.... I mean, to us, communities of practice were small entities: 10, 20, 30, 40 people that really did share a practice...through joint activity. [John Seely Brown]

The argument that Lave and Wenger make is that communities are formed around tensions of membership, continuity, and displacement and it seemed to me that most of those tensions couldn't play out in distal communities where people never came into contact with one another... I tended to see communities of practice as for the most part then face-to-face. [Paul Duguid]

This emphasis on knowledge management of tying technological solutions to "soft" problems is also responsible for the loss of analytic focus referred to in Section 11.3. Coupling CoPs with web technologies in particular coincides with a looser flavor of the theory:

Well, some people call community of practice a collaborative space on the web, you know. So to me that's not a very interesting use of the term. "We have opened a community of practice" means "we have opened a little collaborative space on the web." [Etienne Wenger]

I know there has been an effort to create, for example, web-enabled groups called communities of practice based on people having similar job titles or similar job responsibilities. [Charlotte Linde]

So-called Web 2.0 technologies have been readily labeled by KM practitioners as KM tools (Wilensky et al. 2009). Etienne Wenger regards them as an ideal fit for cultivating CoPs:

This notion of connecting people to people, of enabling group formation not by geography or organizational affiliation but by mutual interest [is done by]... Web 1.0 already but even more [by] Web 2.0 because it puts people really in charge of creating the groups that they want to belong to. It's completely aligned with the notion of communities of practice.

Finally, since IRL was aware that they could not financially depend on Xerox in the long-term, KM provided another potential source for funding:

Actually, funding began to dwindle, down to less and less. It actually stopped... And that's when communities of practice became a part of [the] whole

Doing Business with Theory: Communities of Practice in KM

knowledge management group, etc. *That is how major funds were secured.* [Gitti Jordan, emphasis added]

Once Xerox ceased supporting IRL financially, IRL was always in a precarious funding situation. In fact, we were told by nearly all our informants, when asked why IRL had to shut down, that the reason was financial. Even with outlets such as KM, it was not enough to ensure IRL's survival. Nevertheless, CoP provided a valuable "technology" that KM readily appropriated; this secured KM's modernity.

12. Ethical concerns of theory

Theories are often seen as amoral models that merely seek to reflect the behavior of the real world. The view is that new theories replace old ones because they are more accurate. They come closer to perfectly representing reality. However, communities of practice have unexpected ethical ramifications.

As noted by both Cox (2005) and Osterlund and Carlile (2005), the concept of CoP quickly became a tool targeted toward those in management. Duguid and Lave find this indicative of the domains in which CoP have largely been popularized:

[O]ne of my most severe criticism of the organizational literature is that for "organization" you almost always expect it to read "corporation." And for "organizational member," you're almost always expect it to read "management." [Paul Duguid]

I know it [CoP] got disseminated through Xerox Corporation as a basis for one more fix for the organizational problems of corporate management and then from there to others. [Jean Lave]

Thus, CoP potentially legitimizes a pathway for management to control even the traditionally informal activities that are part of learning. It promises to allow a way to control CoP:

I've said in a couple of places, I think it [knowledge management] was very instrumental. *It was very designed for management. It actually paid remarkably little heed to what I [and Julian Orr] felt it was trying to say about the technicians...* And, I think it seemed to say that knowledge was this... extractable thing that could be managed. And I just think the very concept of knowledge management strikes me as very strange. [Paul Duguid]

This suggests a peculiar way in which CoPs are used. They expand the domain in which management may exercise control. While acknowledging it impossible to police CoPs, Jean Lave expressed strong feelings about such usage:

To the extent that it gets twisted around so that it becomes just another slogan, or another way to control and exploit workers, I'm absolutely opposed to it. I

don't want my work to be used in that way. But I have had to learn that once your work is in the public domain, you lose control of its lives and fortunes. [Jean Lave]

In contrast, CoP has also increasingly been seen as a reason to encourage the proliferation of cross-cutting, hierarchically flattening, heterogeneous groups of people (cf. Section 5.1). In such a case, CoP can empower those who work in the dredges. Here, Lave makes a point to distinguish between “management work” and “business school works”:

What's really in there is the notion that people, in their relations with each other, can engage in serious, creative and complex work. I sometimes think that hierarchical models of work organization do not treat such loci as a condition for the production of complex work—a condition that makes it appealing to the people who do the work. As opposed to the business school works, those engaged in management work (also in e.g., adult education) point to the *Situated Learning* book to confirm that it is legitimate to work as they do. They are using the book in the way I wished it would be used. [Jean Lave]

Susan Stucky also notes that for Hewlett-Packard, CoP became a way to “cross silos” and “cross colonies.” Here we see that *Situated Learning* would empower those who have been undeservedly marginalized by conventional, hierarchical structures of business.

While theories or concepts may be incorrectly used, it is perhaps OK. It is perhaps acceptable if what it does is to achieve a “greater good”:

I've heard active, middle managers comment that the value to them of a notion of “situated learning” or “situated practice” lay in its support the legitimacy of more sideways and less hierarchical management of their activities. I think they have sometimes been able to use notions like “communities of practice” as a kind of ideology for more democratic organization of heterogeneous projects. That is one use of these ideas that appeals to me. [Jean Lave]

Interestingly, while conceptual notions may be used incorrectly, if the end result is humanitarian, then such usage may be acceptable. Wenger describes the subversive, counterculture connotations that came with CoP because it threatened the usual way of doing business:

[T]here was something a bit almost subversive about the notion of the community of practice in an organization: this kind of self-organizing group of people who were kind of taking charge of their own learning [and] who were developing up a practice that was different from the canonical processes of way of doing things.

Doing Business with Theory: Communities of Practice in KM

Finally, Wenger himself acknowledges that some may scoff at the direction CoP has gone simply because it has become a way to make a living:

I know some people feel like I've betrayed the field...[Y]eah, there are people who think that I've kind of like gone in to the dark side...[I] have kind of started to make a living by the notion of cultivating them [CoPs]...[I]t's not just business because I work with government...I work with...school districts. I work...with health care also.

Of interest here is that Wenger purposefully make a point that his work does not merely touch businesses, but also expands beyond to other, more morally redeemable enterprises such as health care and school districts.

Yet, one might ask why making money off a theory is any more a move toward the “dark side” in contrast to making tenure, getting fame, creating disciplines, and so forth in academia? Trying to make a living through a theory's two facets—analytical (descriptive) and prescriptive—is difficult.

I'll tell you, it's easier to make a living in the design-side of the world but I try to keep the analytical perspective alive. [Etienne Wenger]

While there may be less stigma on making a profit off of research done in traditional scientific disciplines (e.g., engineering, biology, or computer science), it seems one needs to explicitly legitimize making a profit on abstract, social theories.

13. A history of reification

In Wenger's (1990) unorthodox thesis, the acknowledgments section reveals one who deeply understands the challenges of theory in academia. His words still ring true today, and it is worth reading his first paragraph (p. x):

Perhaps one of the deepest problems of the academic world is its commoditized, centered view of the production of ideas, reflecting in sharp focus broader tendencies in our civilization....[This is] a characteristic I find profoundly disturbing. I call it “centered” because it promotes a competitive sense of achievement that pointedly centers the development of understanding on specific individuals. Ideas are mine, ideas are yours: careers are built on the ownership of ideas, whose own careers become looted battlegrounds. *The dense fabric of infinite mutual influences, serendipitous connections, and interdependences among contraries—forever invisible—is molded into the epiphany of heroes, whose stature grows beyond life at the cost of heroism in life.* This centered view is at odds with the argument of this thesis as well as with the shared understandings, relations, and hopes that have carried me through execution. [emphasis added]

Our paper has been a modest endeavor to unveil this “dense fabric” of mutual influences that envelop communities of practice. Wenger was acutely aware of the dominant view of science (Hilgartner 1990)—that of a great idea conceived by individuals and heroically disseminated. Our analysis is sympathetic to those who reject such a simplistic view of science, knowledge, and, indeed, history. In fact, part of the motivation for this work is to rediscover history that was marginalized, made invisible.

We believe communities of practice provide a fascinating case study of how abstract ideas and concepts—intangibles—are reified. Technology has often been the focus when one talks of “making something reality.” In his famous case study of Aramis—a failed technological project to provide a publicly guided transport system for Parisians—Latour (1996) describes the relation between *text* and *object* (p. 24):

Aramis was a text; it came close to becoming, it nearly became, it might have become, an object, an institution, a means of transportation in Paris. In the archives, it runs back into a text, a technological fiction...only a fiction that *gains or loses reality* can do justice to the engineers, those great despised figures of culture and history.

Engineers, those tasked to create reality from ideas, are part of this narrative of reification. What about theorists?

As we have hoped to show, theories themselves are subject to the same sociological methods of analysis as technologies. Our intent in this paper is to utilize communities of practice as an exemplar of how ideas, concepts, and theories are translated into more concrete instantiations by actors allied toward particular motivations. While technologies and theories may seem to be on opposite ends of the material spectrum, we found it useful to think of their similarities. In fact, it is a strategic maneuver of actors to socially construct CoP into something more “real”; it then becomes something to be talked about, something to be marketed, something to count, something to own, something to be labeled virtuous, something to evaluate, something to be an expert in, and so on. Just as technological projects such as Aramis (Latour 1996) may start off as scribbled diagrams on napkins and turn into actual prototypes we ride on, theories themselves may start off as amorphous ideas that turn into theories we can point to, promulgate, and assemble.

At any time, CoPs were an “accessory concept” to prompt questions that the burgeoning theory of situated learning might answer—a reaction to the cognitive underpinnings of artificial intelligence, a legitimization for “research” labs (i.e., Xerox PARC), a savior for education, an object passed on to one’s academic progeny, a product to sell, a “fix” for the transmission-model of knowledge in the KM discipline, a formal set of processes to be implemented to foster cross organizational knowledge sharing, a manifesto to “empower” those who did the “actual” work of corporations, and the killer

Doing Business with Theory: Communities of Practice in KM

application for KM. CoP has become what Latour (1996, p. 48) described as a “good swap shop for goals”:

It is appropriate for different groups with divergent interests to conspire with a certain amount of vagueness on a project that they take to be a common one, a project that then constitutes a good “agency of translation,” a good swap shop for goals.

CoP’s attributes and circumstances made it readily transmogrified into more “solid” forms that satisfied the goals of different groups.

From our field work spanning several years, we believe very few people know of the Institute for Research on Learning. Many know of its famous parent, Xerox PARC, but little know (even amongst academics) from whence communities of practice came. The immense popularity of communities of practice is a double-edged sword. Its popularity allowed it to be a “success,” while, at the same time, the popularity of communities of practice obscured its origins.

The organizational form of an institution allowed IRL to become an effective space for creating alliances. It was a key link in the “chain of translations” (Latour 1996, p. 33) that led to the extreme formalization that characterizes CoP as conceived by KM practitioners in the aerospace industry. With all the trappings of a business, IRL deftly packaged, marketed, and sold CoP. IRL was the *obligatory passage point* that promised to cure the *epidemic* affecting businesses—lost knowledge. These translations put a very academically oriented object into the business world, and thus in the hands of those who were inclined to theorize and apply theory in business—knowledge management practitioners.

For KM practitioners in the aerospace industry, CoP has reached a relatively stable form. It is a thing radically transformed from its original conception. As premiere organizational forms of grouping people, CoPs are now tasked to create quality wisdom. They bridge disparate, marginalized, and untapped groups together to collaborate. CoPs provide a blueprint for effective knowledge sharing that takes “people” into account. Specific roles for “members” of a CoP ensure that the “health” of a community remain in check. The formalized process of building a CoP is tightly integrated with technological artifacts deployed by knowledge management experts.

We wish to emphasize that while our inquiry into the story of communities of practice in knowledge management is a significant step toward understanding how ideas, concepts, theories, and other sorts of intangibles traverse spheres of influence, it is neither the *complete* nor *final* story. As we ourselves are subject to the same impartial, subjective, and emotive whims of the actors we have investigated, we cannot claim that how we structure and present our narrative is the truth, nor the best.

Specifically, we do not claim to explain a direct *causal* link between the happenings of IRL and KM (or, more specifically, KM in the aerospace industry).

Our analysis was centered about KM practitioners in the aerospace industry and the transformation of CoP by IRL. For instance, a more comprehensive telling of CoP—one out of scope for this paper—might include an ethnographic inquiry into the various consulting firms (e.g., APQC) who have heavily invested in communities of practice and knowledge management.

We believe this study provides a unique view on an academic theory that was fortuitously translated from academia to business. Its popularity amongst a diverse set of groups is unprecedented, and it is our hope that we have taken a step toward understanding how this came to be. With each translation, this concept, theory, process, or technology succeeded in enrolling interested actors. Some of the translations were slight, some were significant. Nonetheless, CoPs were not simply diffused or transmitted into new worlds. There is “no transportation without transformation” (Latour 1996, p. 119).

14. Conclusion

Although community of practice was initially a reactionary device against learning theories grounded in the detachment of mind from body, its current instantiation ironically harkens back to what it was originally mobilizing against. That there is now a business of communities of practice is indicative of how alluring formalization/reification remains as a way to conceive and manage the world. Despite deliberate efforts to move away from “rationalist, individualist, empiricist conceptions of learning, knowing, and the social world” (Lave 1993, p. 29), one always seems to eventually end up back where he or she started.

Lave (1993) noted that traditional notions of education often divide learning into contextualized and decontextualized forms. And, indeed, all forms of social practice (not just learning) are often viewed through this lens of dualism. This dualistic perspective gives rise to a formalist perspective (Lave 1993, p. 22–23):

[I]t should not be surprising that the head is often also conceived to be a container...for knowledge, while more general knowledge is the container of more particular knowledge, and language is an inert container for the transmission of meaning....[This] should remind us that a formalist view of context is a key conception in conventional theories of action, thinking, knowing, and learning, with, significantly, deep roots much more generally in a Euro-American worldview.

Although situated action methods criticize this as overly simplistic, Lave was right to point out that contextualization and decontextualization *themselves* constitute a practice worthy of analytic inspection. This perspective “leads a robust existence *in practice*, in the contemporary world” (Lave 1993, p. 24). With this line of inquiry, one can ask how is “the myth of an objective, disinterested, asocial world, *made* part of our social practices, in situated ways?” (p. 24).

Doing Business with Theory: Communities of Practice in KM

Thus, it is perhaps doubly ironic that we might be partially answering such a question by investigating the very theory that attempts to mythologize in the first place such an asocial world. That CoP is itself a victim of what it was seeking to attack tells us just how robust and strong such a world view is.

The dominant instantiation of CoP in the aerospace industry is antithetical to Lave's intention for CoP and more importantly, situated learning. It distorted the concept, according to Lave:

The notion of “community of practice” as a naturally occurring object in the world, one you could go out to look for, that is precisely the kind of *transformation of an analytic idea into a natural object* that I mean when I say that *positivist common sense theory simply assimilates and distorts the original intentions of the concept.* [emphasis added]

It was never her goal to create a business of discovering CoPs. To reify CoP meant retrogradely relying on the old, positivist connotations of artificial intelligence (AI) that she was rallying against in *Situated Learning* (Lave and Wenger 1991).

Certainly, we do not mean to imply that the actors we encountered were “cultural dopes” (Garfinkel 1967, p. 245). For example, many researchers were aware of the loss of analytic precision to gain wider acceptance. Some of the KM practitioners in the aerospace industry were also aware that technology itself is not equivalent to CoP. People are cognizant of the disadvantages of the dualist viewpoint on the world. If we were to characterize the path of CoP, it might be said it was full of “compromises” (without any of the negative connotations usually associated with the word). Latour (1996, p. 99–100) asserts that those who fail to compromise will never further their own goals regarding technology.

The only way to increase a project's reality is to compromise, to accept sociotechnological compromises....It's this moral crisis that leads the pure Aramis—the first Aramis, the one that could do everything—to be called nominal, while the series of altered and compromised Aramises is referred to as the *simplified* Aramis, or the *degraded* Aramis, or the *VS* (for *very simplified*) Aramis. If the player is reluctant to compromise his construction, he has lost...The pertinent question is not whether it's a matter of technology or society, but only what is the best sociotechnological compromise.

As the above passage notes, it is an ill-posed question to ask, “Will the real communities of practice stand up?” Just as there are multiple Aramises, there are multiple communities of practice. In other words, the best possible alliance is one of compromise. And, indeed, our actors have compromised intelligently and purposefully with communities of practice.

So, instead, one might regard a community of practice as a concept/theory always on the precipice of this boundary between the belief in overt categorization of the world and belief in the inseparability of such a world into

such categories. Compromises are made to fit this piece of the situated action theory into a situation. In essence, CoP itself is situational.

We offer no final solutions for knowledge management practitioners on how they ought to “use” (if at all) communities of practice. We might envision KM practitioners as the closest embodiment of what a “computer-supported cooperative work (CSCW) practitioner” might be—those concerned with the social and collaborative relation technology has with people—but without the extensive education of an academic. Nonetheless, we believe that if KM practitioners were to gain a wider perspective into the tumultuous history of communities of practice, they would be better equipped (becoming minisociologists, if you will) at balancing between the various instantiations of CoP. The danger, we believe, is when one has descended so far into a particular reality of CoP that it then becomes impossible to see what other realities might offer.

Acknowledgements

Special thanks to Judd Antin, Jeanette Blomberg, Judy Olson, Sally Peters, and Erik Vinkhuyzen for their assistance in the data collection. This work was supported in part by the US National Science Foundation (NSF) under grant 0808783.

References

- Berger, P. L., & Luckmann, T. (1967). *The social construction of reality: A treatise in the sociology of knowledge* (1st ed.). New York: Anchor.
- Brown, J. S., & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*, 2(1), 40–57.
- Brown, J. S., & Duguid, P. (2001). Knowledge and organization: a social-practice perspective. *Organization Science*, 12(2), 198–213.
- Butterfield, H. (1965). *The Whig interpretation of history*. New York: W. W. Norton & Company.
- Callon, M. (1996). Some elements of a sociology of translation: Domestication of the scallops and the fishermen of Saint Brieuc Bay. In J. Law (Ed.), *Power, action and belief: A new sociology of knowledge?* (pp. 67–78). London: Routledge.
- Contu, A., & Willmott, H. (2003). Re-embedding situatedness: the importance of power relations in learning theory. *Organization Science*, 14(3), 283–296.
- Cooper, G., & Bowers, J. (1995). Representing the user: Notes on the disciplinary rhetoric of HCI. In P. J. Thomas (Ed.), *The social and interactional dimensions of human-computer interfaces* (pp. 49–66). Cambridge, UK: Cambridge University Press.
- Cox, A. (2005). What are communities of practice? A comparative review of four seminal works. *Journal of Information Science*, 31(6), 527–540.
- Davenport, T. H. (1994). Saving IT's soul: human-centered information management. *Harvard Business Review*, 72(2), 119–131.
- Dourish, P., Hayes, G. R., Irani, L., Lee, C. P., Lindtner, S., Nardi, B., et al. (2008). Informatics at UC Irvine. In: ACM Extended Abstracts on Human Factors in Computing Systems (CHI'08) (pp. 3651–3656). New York: ACM Press.

Doing Business with Theory: Communities of Practice in KM

- Duguid, P. (2008). Prologue: Community of practice then and now. In A. Amin & J. Roberts (Eds.), *Community, economic creativity, and organization* (pp. 1–10). New York: Oxford University Press.
- Feldman, M. S., & March, J. G. (1981). Information in organizations as signal and symbol. *Administrative Science Quarterly*, 26(2), 171–186.
- Fischer, G. (2001). Communities of Interest: Learning through the interaction of multiple knowledge systems. In: 24th Annual Information Systems Research Seminar In Scandinavia (IRIS'24) (pp. 1–14). Ulvik, Norway.
- Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall.
- Hilgartner, S. (1990). The dominant view of popularization: conceptual problems, political uses. *Social Studies of Science*, 20(3), 519–539.
- Klein, H. K., & Hirschheim, R. (2008). The structure of the IS discipline reconsidered: implications and reflections from a community of practice perspective. *Information and Organization*, 18(4), 280–302.
- Latour, B. (1996). *Aramis, or the love of technology*. Cambridge, MA: Harvard University Press.
- Lave, J. (1988). *Cognition in practice*. Cambridge, UK: Cambridge University Press.
- Lave, J. (1993). The practice of learning. In S. Chaiklin & J. Lave (Eds.), *Understanding practice: Perspectives on activity and context* (pp. 3–32). Cambridge, UK: Cambridge University Press.
- Lave, J. (2008). Epilogue: Situated learning and changing practice. In A. Amin & J. Roberts (Eds.), *Community, economic creativity, and organization* (pp. 283–296). New York: Oxford University Press.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, UK: Cambridge University Press.
- Leonard, D., & Kiron, D. (2002, October). Managing knowledge and learning at NASA and the Jet Propulsion Laboratory (JPL) (Tech. Rep. No. 9-603-062).
- Orr, J. E. (1996). *Talking about machines: An ethnography of a modern job*. Ithaca, NY: Cornell University Press.
- Osterlund, C., & Carlile, P. (2005). Relations in practice: sorting through practice theories on knowledge sharing in complex organizations. *The Information Society*, 21, 91–107.
- Palla, M., & Ward, M. (2006). *Connecting value stream white-spaces through knowledge-based strategies and CoPs*. Las Vegas: Raytheon Bridges the Gaps.
- Phillips, N., & Hardy, C. (2002). *Discourse analysis: Investigating processes of social construction* (2nd ed.). Thousand Oaks, CA: Sage Publications.
- Rumizen, M. C. (2002). *The complete idiot's guide to knowledge management*. Indianapolis, IN: Alpha Books.
- Strauss, A., & Corbin, J. (1998). *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Thousand Oaks, CA: Sage Publications.
- Su, N. M., Wilensky, H., Redmiles, D., & Mark, G. (2007). The gospel of knowledge management in and out of a professional community. In: Proceedings of the ACM Conference on Supporting Group Work (GROUP'07) (pp. 197–206). New York: ACM Press.
- Suchman, L. (1987). In: R. Pea, J. S. Brown, & C. Heath (Eds.). Plans and situated actions: The problem of human-machine communication. Cambridge, UK: Cambridge University Press.
- Sutton, S., Betser, J., Hornickel, M., Gregorio, M., Kern, J., Lincoln, C., et al. (2008). The Aerospace Corporation case study. In J. Liebowitz (Ed.), *Knowledge retention: Strategies and solutions* (1st ed.). Boca Raton, FL: Auerbach.
- The Boeing Company, Rocketdyne Division. (2003). In: P. Leavitt (Ed.). Using knowledge management to drive innovation (Spiral ed., pp. 117–142). Houston, TX, USA: American Productivity & Quality Center.

- Tufte, E. (2003). *The cognitive style of PowerPoint: Pitching out corruptions within*. Chesire, CT: Graphics Press.
- Vann, K., & Bowker, G. (2001). Instrumentalizing the truth of practice. *Social Epistemology*, 15, 247–262.
- Wenger, E. (1990). Toward a theory of cultural transparency: Elements of a social discourse of the visible and the invisible. Unpublished doctoral dissertation, University of California, Irvine.
- Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. New York: Cambridge University Press.
- Wenger, E., McDermott, R. A., & Snyder, W. (2002). *Cultivating communities of practice*. Boston, MA: Harvard Business Press.
- Wilensky, H. N., Su, N. M., Redmiles, D., & Mark, G. (2008). A community of knowledge management practitioners: mirroring power across social worlds. In: Proceedings of the 20th IFIP World Computer Congress, WG12.6 Conference on Knowledge Management in Action (KMIA'08) (pp.195–207). Boston, MA: Springer.
- Wilensky, H. N., Redmiles, D. F., & Su, N. M. (2009). The dissemination of knowledge management. In: Proceedings of the ACM Conference on Supporting Group Work (GROUP'09) (pp. 199–208). New York: ACM Press.
- Williams, R. (1985). *Keywords: A vocabulary of culture and society* (Revised edth ed.). New York: Oxford University Press.
- Wilson, T. D. (2002). The nonsense of 'Knowledge Management'. *Information Research*, 8(1), paper no. 144. (Available at <http://InformationR.net/ir/8-1/paper144.html>).
- Zizek, S. (2008). *The sublime object of ideology* (2nd ed.). London: Verso.