
Toward a Theory of Knowledge Reuse: Types of Knowledge Reuse Situations and Factors in Reuse Success

M. LYNNE MARKUS

M. LYNNE MARKUS is Professor (Chair) of Electronic Business at the City University of Hong Kong. She is on leave from the Peter F. Drucker Graduate School of Management, Claremont Graduate University, where she is Professor of Management and Information Science. Professor Markus's research focuses on electronic commerce, enterprise systems integration, and knowledge management. Dr. Markus was formerly a member of the faculties of the Anderson Graduate School of Management (UCLA) and the Sloan School of Management (MIT). She has also taught at the Information Systems Research Unit, Warwick Business School, UK (as Visiting Fellow), at the Nanyang Business School, Singapore (as Shaw Foundation Professor), and at the Universidade Tecnica de Lisboa, Portugal (as Fulbright/FLAD Chair in Information Systems). Dr. Markus has received research grants and contracts from the National Science Foundation, the Office of Technology Assessment (U.S. Congress), the Advanced Practices Council of SIM International, the Financial Executives Research Foundation, and Baan Institute. She is the author of three books and numerous articles in journals such as *MIS Quarterly*, *Management Science*, *Organization Science*, *Communications of the ACM*, and *Sloan Management Review*. She serves on the editorial boards of several leading journals in the information systems field. She has served as AIS Council member for the Americas and as VP for Academic Community Affairs for SIM International. Markus holds a B.S. in Industrial Engineering from the University of Pittsburgh and a Ph.D. in Organizational Behavior from Case Western Reserve University.

ABSTRACT: This paper represents a step toward a theory of knowledge reusability, with emphasis on knowledge management systems and repositories, often called organizational memory systems. Synthesis of evidence from a wide variety of sources suggests four distinct types of knowledge reuse situations according to the knowledge reuser and the purpose of knowledge reuse. The types involve shared work producers, who produce knowledge they later reuse; shared work practitioners, who reuse each other's knowledge contributions; expertise-seeking novices; and secondary knowledge miners. Each type of knowledge reuser has different requirements for knowledge repositories. Owing to how repositories are created, reusers' requirements often remain unmet. Repositories often require considerable rework to be useful for new reusers, but knowledge producers rarely have the resources and incentives to do a good job of repurposing knowledge. Solutions include careful use of incentives and human and technical intermediaries.

KEY WORDS AND PHRASES: collaboration, communities of practice, experts, group work, intermediaries, knowledge management, knowledge repositories, knowledge reuse, novices, organization memory, teams

ONE OF THE KEY THEMES IN KNOWLEDGE MANAGEMENT today is the role of information technology (IT) in the transfer of knowledge between those who have it and those who don't. It is widely acknowledged that knowledge has two dimensions—*explicit* (knowledge that has at minimum been “captured” and articulated and has ideally been “codified,” that is, documented, structured and disseminated) and *tacit* (knowledge that resides in people's heads or “muscle memory” and may be destined to remain there). Only explicit knowledge is the province of information technology, including the communication systems by which people informally share their observations and the more formal repositories in which structured knowledge is stored for later reuse. Some people hold that knowledge repositories play a relatively unimportant role in knowledge reuse, arguing that face-to-face communication and good knowledge sharing processes between the sources and intended recipients of knowledge are the keys to successful knowledge reuse. Others are much more sanguine about the potential for formal repositories.

The purpose of this paper is to begin building a theory of knowledge reusability, with particular emphasis on the role of knowledge management systems [5] and knowledge repositories, often called *organizational memory systems* [1] or *organization memory information systems* [33]. Although much knowledge reuse involves *access to experts*, not *access to codified expertise* [13, 51, 52], it is increasingly the case that the identification and selection of experts are mediated through knowledge management systems [7, 52]. Therefore, both access to experts and access to expertise are included within the scope of this paper.

The primary source of material for this nascent theory is published accounts of situations involving the creation and use of written and computer-based records for the purposes of preserving, accessing, and reusing knowledge about what was done, how and why things are or were done, what things mean, and how this knowledge can be applied in other settings. In other words, the focus is not on the creation and use of primarily numerical records of the sort we usually call databases, but on the less structured information that is usually stored as documents.

The process of theory development involved categorizing and naming observations and creating linkages among them and developing tentative hypotheses and comparing them against other observations. The ultimate goal of this exercise is to articulate a design theory [70] that specifies the conditions under which successful knowledge reuse is likely to occur. Much work remains to realize this goal. The major contributions of this paper are, first, to identify four types of knowledge reuse situations involving different knowledge “reusers” (shared work producers, shared work practitioners, expertise-seeking novices, and secondary knowledge miners), differentiated by the “knowledge distance” between those who have the knowledge and those

who don't, and second, to outline what needs to be done to make repositories useful for the different types of knowledge reusers.

Background

KNOWLEDGE PROCESSES ARE OFTEN CATEGORIZED by whether they involve *knowledge creation* (as in research or new product development) or *knowledge reuse* (as in sharing best practices or helping others solve common technical problems) [20]. Knowledge creation is often viewed as somehow more important, more difficult to manage, and less amenable to information technology support [50]. However, the effective reuse of knowledge is arguably a more frequent organizational concern and one that is clearly related to organizational effectiveness [22, 58, 64].

Despite its importance, knowledge reuse is something that we know relatively little about. Although knowledge reuse has been observed and researched under many different names in many different settings, findings about knowledge reuse have remained relatively dispersed and unintegrated. One possible explanation is that knowledge reuse is seen as a unitary phenomenon—pretty much the same regardless of who does it, how, and why.

One recent and notable exception is the work of Nancy Dixon, whose book, *Common Knowledge* [22], identified five different types of knowledge transfer situations called serial transfer, near transfer, far transfer, strategic transfer, and expert transfer. For each of these situations, Dixon offered “design guidelines” for successful knowledge management. This very useful design theory has some points of similarity with the arguments in this paper (particularly an emphasis on the recipients of transferred knowledge), but it differs from the one outlined here in several ways.

First, the focus and methodology are different. Dixon's primary concern is with procedural knowledge—knowledge about how to do something better (often called “best practice”). She examined only successful knowledge management projects. She abstracted from her examples the characteristics that differentiate transfer situations and the principles that appeared to make the difference in knowledge transfer success. She is not explicitly concerned with the role of knowledge repositories. In fact, only two of her five types involve repositories, and the implication is that repositories are not useful in the other situations. By contrast, I am concerned with all types of knowledge that people might need to reuse in different situations (e.g., factual, analytic, and rationale knowledge, in addition to procedural). I have examined both successful and unsuccessful situations involving the creation and reuse of knowledge repositories to identify four different types of knowledge reuse situations (involving shared work producers, shared work practitioners, expertise-seeking novices, and secondary knowledge miners). I believe that repositories can play a role in almost all knowledge reuse situations—the trick is to specify the knowledge reuse needs in each situation and the implications for repository design and related interventions.¹

A second difference is the concepts used to differentiate the knowledge transfer or reuse situations. Dixon's key concepts are: (1) recipient of knowledge transfer (is the

recipient also the source?), (2) task (frequency and routineness of work performance), and (3) knowledge (explicit versus tacit). Her work suggests that repositories are successfully used only when the recipient is not the source of knowledge, when the task is routine, and the knowledge is explicit. My concepts are: (1) knowledge reuser and the purpose of knowledge reuse, (2) what the recipient needs to know, knows, and doesn't know, and (3) challenges the recipient faces at each stage of knowledge reuse (defining the question, locating experts or documents, selecting experts or documents, and applying the knowledge). In both cases, our dependent variable is the same—successful knowledge transfer or reuse.

The next section of this paper provides an overview of basic concepts in knowledge reuse. Subsequent sections present a typology of different knowledge reuse situations, discuss the nature of repositories used by the different types of knowledge reusers, and outline some additional elements in successful knowledge reuse.

Basic Concepts in Knowledge Reuse

MEANINGFUL DISCUSSION OF KNOWLEDGE REUSE requires consensus on basic definitions. In this section, the process of knowledge reuse, the roles in knowledge reuse, and the types of knowledge repositories are defined and briefly described.

Process

The knowledge reuse process can be described in terms of the following stages: *capturing or documenting* knowledge, *packaging* knowledge for reuse, *distributing or disseminating* knowledge (providing people with access to it), and *reusing knowledge* [4, 38].

Capturing and documenting knowledge can occur in at least four ways. First, documenting can be a largely passive by-product of the work process, as when virtual teams or communities of practice automatically generate archives of their informal electronic communication that can later be searched [3, 42, 48, 62]. Second, documenting knowledge for potential reuse can occur within a structure such as that provided by facilitators using brainstorming techniques, perhaps mediated by the use of electronic meeting systems [21], or using a technique like gIBIS analysis, perhaps mediated by design rationale capture systems [16]. Third, documenting can involve creating (pre)structured records (for example, of technical support interventions [24, 61]) as part of a deliberate, before-the-fact knowledge reuse strategy. And, fourth, documentation can involve a deliberate, after-the-fact strategy of filtering, indexing, packaging, and sanitizing knowledge for later reuse, as in the creation of learning histories [65], consultants' "Power Packs" [39], expert help files [1], or the creation of a data warehouse [10].

Packaging knowledge is the process of culling, cleaning and polishing, structuring, formatting, or indexing documents against a classification scheme. Among the activities involved in knowledge packaging are: authoring knowledge content [65], codify-

ing knowledge into “knowledge objects” by adding context [65], developing local knowledge into “boundary objects” by deleting context [1, 2], filtering and pruning content [26], and developing classification schemes [47].

Distributing knowledge can be passive, such as publishing a newsletter or populating a repository for users to browse, or active, such as convening an “After Action Review” meeting or “pushing” knowledge via an electronic alert to those who need to know [22]. Also included in the category of knowledge dissemination are a variety of *facilitation* activities, such as assessing knowledge reuse needs [67], helping the intended users use knowledge or knowledge management tools in appropriate ways [37], helping organizations understand the need to adopt newly codified best practices [13], and facilitating the development of internal or external communities [14, 23].

Reusing knowledge, the final phase, has been said to involve both *recall* (that information has been stored, in what location, under what index or classification scheme) and *recognition* (that the information meets the users’ needs [41], as well as actually applying the knowledge). Similarly, the use of human expertise involves both the *identification* of experts in a subject matter and the *selection* of the expert most appropriate for a particular query [52], as well as query, response, and application of the results. An important type of knowledge reuse involves the systematic secondary analysis of records created for very different purposes, which is often called *data mining* [11]. For the purposes of this paper, the reusing knowledge phase consists of four different activities. The first is *defining the search question*. This step is essential for successful reuse. It has often been noted that one characteristic separating experts from novices is that experts know what questions to ask. The second is *the search for, and location of, experts or expertise*. The third is *selection of an appropriate expert or of expert advice* from the results of the search. The last is *applying the knowledge*, which may involve analysis of general principles against a specific situation—a process sometimes called “recontextualization” of knowledge that was decontextualized when it was captured and codified [1, 2, 12, 41].

Roles

There are three major roles in the knowledge reuse process: *knowledge producer*—the originator and documenter of knowledge, who records explicit knowledge or makes tacit knowledge explicit, *knowledge intermediary*—who prepares knowledge for reuse by eliciting it, indexing it, summarizing it, sanitizing it, packaging it, and who performs various roles in dissemination and facilitation, and *knowledge consumer*—the knowledge reuser, who retrieves the knowledge content and applies it in some way.

This observation by itself does not get us very far, until it is noted that the three roles can be performed by the same individual(s) or group(s), different individual(s) or group(s), or some combination. For example, at the end of each semester I (try to) tidy up my paper and electronic course files so that I’ll have an easier time next semester. In this example, I am the producer/documenter of (some of) the explicit knowledge in my course files, the packager/intermediary, and the intended reuser.

(People who document knowledge for their own reuse are sometimes called “prosumers” [32].) In the case of Andersen Consulting’s packaged knowledge products (intended for sale to, and use by, clients), the producers of the knowledge are Andersen consultants, the packagers are knowledge intermediaries employed by the firm to make the knowledge reusable, and the intended users are Andersen clients.

It is also possible that some of these roles may be performed by information technology. Information technology affords particular opportunities as an intermediary in knowledge reuse by automatically categorizing, abstracting, filtering, and disseminating documents. In the case of some systems based on artificial intelligence, parts of the knowledge producer and consumer roles may also be performed by technology.

Types of Repositories

Many kinds of repositories are involved in knowledge reuse. Perhaps the most basic distinction is that between repositories of *documents* and repositories of *data* [12]. Blair observed that retrieving information in the form of textual documents is fundamentally different than retrieving data. Consequently, strategies for indexing and storing the different kinds of information must differ as well [12]. Increasingly, this basic distinction must be augmented with graphical information such as engineering drawings [42] and with audio, video, and multimedia “documents.”

Davenport and colleagues [19] distinguish among repositories that store *external* knowledge (e.g., demographic data, competitive intelligence), those that store structured *internal* knowledge (in the form of data or documents), and those that store *informal* information (such as transcripts of group discussions via electronic meeting systems, computer conferencing, or electronic mail). Alavi and Leidner [5] mention several specific kinds of (semi)structured internal knowledge repositories, including corporate yellow pages and people information archives.

A further way of classifying repositories is by means of the kind of knowledge they contain. Zack [74] contrasts *general knowledge* (including, for example, explicit scientific knowledge) and *specific knowledge* (including knowledge of the local context). (See Choudhury and Sabherwal [15] for a similar distinction between technical and contextual knowledge.) Moorman and Miner [55] distinguish between *declarative knowledge* (knowledge about facts) and *procedural knowledge* (knowledge about how things are done). To this distinction must be added *rationale knowledge*—knowledge about why things were done, as in the design of software products [56]—and *analytic knowledge*—the conclusions reached by applying declarative and procedural knowledge to a particular fact domain [27].

A Typology of Knowledge Reuse Situations

REVIEW OF ACADEMIC AND PRACTICAL WRITINGS on the use of documents and knowledge repositories suggests that there are at least four different types of situations in

tions in which knowledge is reused. The basic characteristics differentiating these types are the knowledge reuser (in relation to the source or producer of knowledge) and the purposes of reuse. The four types are: reuse by shared knowledge producers, reuse by shared work practitioners, reuse by expertise-seeking novices, and reuse by secondary knowledge miners. The typology is summarized in Table 1.

Shared Work Producers

Knowledge reusers can be close to or distant from those who produced the knowledge, where distance is measured in terms of shared knowledge. Closest are those who actually produced the knowledge themselves while working on a *shared work product*. There are at least two types of shared work producers: *homogeneous work groups* (e.g., software development and support teams and physicians in a medical practice) and *cross-functional teams* (e.g., new product development teams and consulting project teams).

Examples of knowledge reuse in homogeneous work groups include software support workers accessing the history of prior diagnostic and repair details [24, 61], software development and ERP system implementation team members revisiting design decisions later in the project [9, 16], and physicians reviewing prior entries in patients' medical records [35]. Examples of reuse in cross-functional work teams include members of a high tech product design team evaluating proposed design solutions [42, 44], members of an invention assessment team providing feedback to the inventors following an assessment panel meeting [57], and members of a consulting team modifying other teams' work products when preparing a proposal for a new client.

Because shared work producers create and document the knowledge they later reuse, they have fewer challenges in knowledge reuse than do other types of reusers. However, they do have some challenges, such as sketchy note taking, bad filing systems [45], and faulty memory [57]. Even when they keep good records, they may have trouble finding what they are looking for. Electronic records produced as a by-product of the work itself, such as consultants' work papers or engineering drawings [43], can quickly become so voluminous that members find them cumbersome to search. For instance, one virtual team committed to using a sophisticated knowledge management system found that they could easily spend 10 minutes out of a 45-minute team meeting searching a 1,000-entry knowledge base for the information they needed [44]. These problems were so severe that team members advocated the use of knowledge intermediaries to help them cope: "The [aerospace new product design] team never did resolve the issues of speedy knowledge retrieval. In the end, they recommended that [virtual] teams should consider establishing a role of a Knowledge Manager" [44, pp. 15–16].

Once individuals or groups have located their own previously stored relevant explicit knowledge, they have relatively little difficulty deciding what information is useful and putting the knowledge to effective use, because they generally "understand one another's implicit knowledge and assumptions" [1, p. 7]. Therefore, they

Table 1. Types of Knowledge Reuse Situations

	Shared Work Producers	Shared Work Practitioners	Expertise-Seeking Novices	Secondary Knowledge Miners
Description	People working together on a team, either homogeneous or cross-functional; producers of knowledge for their own later reuse	People doing similar work in different settings; producers of knowledge for each other's use	People with an occasional need for expert knowledge that they do not possess and do not need to acquire themselves because they need it rarely	People who seek to answer new questions or develop new knowledge through analysis of records produced by other people for different purposes
Purpose of Knowledge Reuse	<ul style="list-style-type: none"> • Keep track of current status and things needing attention • Recall reasons for decisions when decisions need to be revisited or when there is turnover among team members • Learn how the team can perform better on the next project 	<ul style="list-style-type: none"> • Acquire new knowledge that others have generated (e.g., how to handle a particular type of problem) • Get advice about how to handle a particularly challenging or unusual situation that is new to the team • Gain access to observations that spur innovations 	<ul style="list-style-type: none"> • Answer an arcane question or solve an ad hoc problem • Approximate the performance of experts • Minimize the need for experts 	<ul style="list-style-type: none"> • Seek answers to new questions or create new knowledge

What Reusers Need to Know, Know, and Don't Know

- Users need to know what was done (declarative, factual knowledge), how it is/was done (procedural knowledge), why it was done (rationale), what could be done better (analytic knowledge)
- Users share general knowledge and specific/contextual knowledge related to project and their area of expertise/involvement
- Users may not know about the work of other team members, particularly those in cross-functional teams
- Users need to know how to do something and why a particular procedure works
- Users share general knowledge and share knowledge of what contextual knowledge is useful
- Users may not have specific contextual knowledge of the producers' settings
- Users need access to others' expertise (e.g., an answer to a question) without actually needing to acquire the expertise itself (e.g., the ability to answer related questions)
- Users have knowledge of the local context
- Users do not have relevant general knowledge
- Users do not know what aspects of local context are important
- Users may not know how to analyze general knowledge against specific context knowledge
- Users need solutions for novel problems
- Users have general analytic expertise
- Users often lack the general and specific knowledge of the people creating the records

Challenges Reusers Experience (and Strategies They Use) When

- Defining the Search Question
 - Minimal problem in homogeneous shared work teams since members share general and specific knowledge
 - More challenging in cross-functional teams
- May not know they need expert advice
- May lack knowledge of expert jargon
- May not be able to recognize technical "symptoms" in local context
- May be unable to articulate the question or problem well
- Defining the question will be especially challenging in the case of knowledge discovery

(continued)

Table 1. Types of Knowledge Reuse Situations (Continued)

	Shared Work Producers	Shared Work Practitioners	Expertise-Seeking Novices	Secondary Knowledge Miners
• Locating Experts or Knowledge/Expertise	<ul style="list-style-type: none"> • Teams frequently keep good records about what they did as a by-product of the work, but they often forget the rationales for their decisions after some period of time has elapsed; problem is compounded by turnover • Teams often experience difficulty locating the information they need in work "transcripts" 	<ul style="list-style-type: none"> • Practitioners use networks of contacts to locate experts/expertise 	<ul style="list-style-type: none"> • May have great difficulty locating suitable experts because of difficulty defining the problem 	<ul style="list-style-type: none"> • May have difficulty identifying repositories likely to contain useful information • May have difficulty finding or creating appropriate search or discovery algorithms
• Selecting Experts or Expertise	<ul style="list-style-type: none"> • Not usually a problem 	<ul style="list-style-type: none"> • Use knowledge of reputations to assess quality of experts/expertise 	<ul style="list-style-type: none"> • Lack suitable criteria for judging quality of experts/expertise 	<ul style="list-style-type: none"> • "Spurious results" are a common problem • Results should be triangulated and pilot tested
• Applying the Knowledge	<ul style="list-style-type: none"> • If they can find what they are looking for, this is not usually a problem 	<ul style="list-style-type: none"> • Usually have little difficulty applying the expertise, once it has been selected 	<ul style="list-style-type: none"> • May lack ability to apply good answers/advice successfully 	<ul style="list-style-type: none"> • Not usually a problem

<p>Recommendations for Promoting Successful Reuse</p>	<ul style="list-style-type: none"> • Maintain context in the record • Provide support for indexing and searching (e.g., periodically summarize transcript threads and purge old records) • Require documentation of rationale knowledge • Do not provide public access to these repositories 	<ul style="list-style-type: none"> • Repackage knowledge, providing quality assurances (e.g., authorship), freshness dating, and appropriate indexing and searching capabilities • Decontextualize knowledge, but publish context information along with the content • Provide access to experts as well as to packaged expertise • Push packaged knowledge to appropriate recipients • Provide appropriate incentives for contributions and reuse 	<ul style="list-style-type: none"> • Repackage knowledge, decontextualizing it, but provide support for recontextualization in the local context • Make heroic efforts to translate knowledge into terminology that novices can understand and search • Provide access to experts as well as expertise • Provide awareness training and consultation • See Markus et al. [50] for some additional suggestions 	<ul style="list-style-type: none"> • Store context information (i.e., metadata) with all repositories to facilitate secondary reuse • Provide thorough training in knowledge base structures • Provide thorough training in analysis, synthesis and drawing valid conclusions • Verify all results (e.g., conduct pilot tests)
---	--	---	--	--

can more easily understand and deal with contextual information in the documentation that might be “incorrect, incomplete, or incoherent,” and they can successfully reuse the raw, unprocessed records that are created as a by-product of knowledge work. For example, Finnish paper mill workers used their electronic diary, described as a form of “talking out loud,” to catch up with what had happened during their periods away from the mill: “It has now become normal for workers returning from [5] days off to ‘catch up’ on those days from the Diary. This did not happen with the paper diary” [40, p. 54].

Once they have found what they are looking for, not only can shared work producers reuse their own raw records, they may actually find the documentation difficult to reuse if information about the context has been stripped away [31], as often happens when the work record has been sanitized and processed by intermediaries. For an example of the difficulties here, see the appendix entitled “An Example of Reuse Requiring Contextualized Information.”

Shared Work Practitioners

A second type of knowledge reusers is quite different from shared knowledge producers in terms of the structure of their communications and the archives they produce [73]. Shared knowledge *practitioners*, people who share a *community of practice* [71, 72], include specialists who occupy the same roles in different locations, work units, or organizations, such as consultants in a practice area² [6, 34], oil field maintenance workers [37], and human resource management professionals [18]. Shared work practitioners produce knowledge for *each other* to use.

Even though the members of a community of practice share general knowledge and knowledge about the important dimensions of context, they may have considerable difficulty reusing knowledge produced by other members of their community. First, knowledge workers may have difficulty *selecting* from the available documents those that are most appropriate to their needs [12, 41]. They may be unsure, for example, whether knowledge documented in repositories is current or out-of-date [26]. Because of such difficulties, consultants at Andersen Consulting often asked their colleagues about which documents in the firm’s repositories would be most helpful: “To avoid situations [of overload and retrieving poor quality documents], consultants often called colleagues who they knew to ask what documents were good for particular applications. This referral system was based on interpersonal trust” [38, p. 9].

In addition, the reputation of the people who had contributed documents was important in document selection: “Some [professionals’] contributions . . . were often used by others” [38, p. 9]. Similarly, Orlikowski [61] observed that the technical support personnel at Zeta used authorship of entries to assess the quality of database entries:

An interesting metric developed by the specialists to assess data quality was their use of incident authorship as an indicator of quality. Each incident that is entered is automatically assigned a unique number, which includes a code identifying the particular specialist who entered it. . . . You tend to evaluate information

differently from different people. So if you see 40 items from a search you go to the incidents of those folks you've gotten good information from in the past. . . . I know that Arthur has a reputation for writing shorts novels as resolutions. I mean, he's a wonderful source of information . . . So when I get an incident from him, I'm very comfortable with that information. Whereas, some of the other people in the department will put in one or two sentence resolutions. And it tends to make it a little vaguer and more difficult to be confident about. [61, p. 12]

The identification of a record's author also plays an important role in interpreting medical records: "[I]t has long been argued that the doctors' ability to recognise the handwriting of their colleagues, and therefore who saw which patient for what, is an invaluable resource for making sense of the consultation" [35, p. 359].

Finally, shared work practitioners may lack the contextual knowledge needed to interpret the documents [1]. Whereas this problem is most acute when the consumers are very different from producers (e.g., novices vs. experts), it can even be a problem for "competent" members of the same community of practice. For example, describing the challenges of accessing paper documents (prior to the electronic knowledge base) in a consulting firm, Galunic and Weeks [26] wrote:

Each partner kept documents relating to the assignments he or she had worked on in personal filing systems . . . a cupboard which had photocopies of some of the documents that had been done. . . . It was hard to find documents relating to a particular industry or approach, and anyway the most relevant documents were not publicly available. This was for two reasons. The first was client confidentiality. . . . The second problem was that even when a consultant found the right person and convinced him or her to share documents, the documents could not be understood without having been part of the assignment. Copies of the presentations given to clients . . . were too context-specific to be of general use. [26, p. 3]

Similar points have been made about the documents stored in the electronic knowledge repositories of consulting firms [38].

Because of the many problems shared work practitioners have in finding relevant and high quality knowledge for reuse, Dixon [22] recommends the use of intermediaries to "push" filtered and packaged knowledge to potential reusers.

Expertise-Seeking Novices

In addition to situations in which knowledge producers and consumers are the same or similar, many knowledge reuse projects are intended to support people who differ substantially from the knowledge creators. This situation has been called *knowledge transfer* (as opposed to knowledge sharing with colleagues), and the situation often involves novices' access to experts and expertise. Examples include manufacturing workers using an expert system for organization design knowledge [50], consulting firms' clients making use of structured expertise as in ASK Ernie [39], and customers accessing a firm's technical support FAQ lists.

The more dissimilar the reusers are from the knowledge generators in terms of knowledge, the more difficulty they may have in defining the search question, locating and selecting experts and expertise, and reusing even carefully packaged knowledge. First of all, they may not know the jargon, the right questions to ask, or the right “symptoms” to report. Second, they usually require information that has been carefully “decontextualized”: “Otherwise, the secondary users will drown in unnecessary, unhelpful, or conflicting data” [2, p. 46]. Third, they require information to be presented to them in a very accessible way, whether this involves technology or a human “high-touch” interface.

For example, El Sawy and colleagues [24] described an expert system that had successfully employed case-based reasoning to support the knowledge reuse needs of customer support personnel in a high tech company. But when they attempted to extend the use of this system to their product resellers, they were unsuccessful. The resellers did, however, seek information directly from people at the company via e-mail:

In 1994, Storage Dimensions tried to give its resellers direct access to TechConnect from their remote computers . . . with fully GUI features. . . [The resellers] never used it. Apparently, for the casual user trying to play the role of technical support engineer, the full functionality and richness features of TechConnect were beyond what a casual user was willing to remember. On the other hand, the TechConnect e-mail and internet connection are very successful. [24, p. 479]

Interestingly, Storage Dimensions was able to provide direct access to some of TechConnect’s features to customers, who as “technical support engineers” are arguably even more “causal” than are the company’s resellers. But the capability successfully used by customers was not the “full functionality and richness” used by internal support personnel and offered to the resellers, but rather a stripped-down version:

[T]he TechConnect web access route allows a customer to submit problem symptoms to TechConnect that will then go search its knowledge base, make some computations that go beyond key-word search, and return with a list of problem solution documents. . . [But] currently TechConnect self-help does not allow direct knowledge base access. [24, pp. 478–479].

In addition, the stripped down knowledge capability provided to customers was augmented by a good high-touch customer support system: “Thus self-help should only be [offered to customers] after a support staff is in place” [24, p. 479].

Not surprisingly, other organizations that have successfully packaged knowledge for reuse by customers and other novices have either gone to extraordinary lengths to make the knowledge easily accessible [25, 50], or they have provided a very “high-touch” service that involves the intermediation of human experts, as in the case of Ask Ernie [68].

Brøderbund went the route of making its customer self-help very easy to use:

[T]he key is knowing who your audience is. I look at this as we have two separate audiences. We have an internally trained support staff, and then we have the customer out there. I don't expect the customer to be trained, so we have to build case bases that they find easy to understand and use. The internal site is very sanitized, and you have to know a little bit about the systems to use it. The Web site is very intuitive; it's based on a lot of graphical clicking, so even a young kid could solve their problem on the Internet. [25]

Ernst & Young, on the other hand, went the route of providing their smaller customers access to packaged knowledge through the helpful intermediation of expert consultants:

Simply offering online access to a database full of documents would not be much of a help to clients, however. . . . when clients access the Ernie Web site, they are asked to formulate a question, assign a title to it, and offer some background on how they plan to use the information. Then they choose one of eight categories . . . and submit the question. . . . Once the query is formulated and submitted, Ernie automatically routes it to the appropriate department—and the appropriate consultant within the department. . . . The consultant is then responsible for packaging together a response that includes his or her specific experience in the subject area, as well as data from Ernst & Young's resource database. [68]

Secondary Knowledge Miners

Perhaps the most extreme case of reuse involves data mining, in which analysts attempt to extract knowledge from records that were collected by others, possibly unknown to the reuser, for very different purposes. Studies show that data mining is most successfully practiced by highly trained analysts who have extensively studied the structures and limitations of their datasets and been coached in the problems involved in drawing inferences from secondary data analysis. For example, Bashein and colleagues found that data miners at BankAmerica received six months of training and apprenticeship before they were allowed solo access to the company's data mining facilities [10]. Although most research on data mining has focused on structured data, similar issues are likely to apply in the case of secondary reusers of documents.

Summary

Table 1 outlines a typology of knowledge reuse situations organized by the type of knowledge reuser. Shared work producers are creators of the knowledge they later consume. Consequently, they have relatively few problems with knowledge reuse. Their primary problems are capturing appropriate information (especially about design rationale) and searching through the record to find what they need. Shared work practitioners are members of a community of practice who reuse knowledge produced by other members of the community, so they have few difficulties applying knowledge once they have located and selected it. But location and selection are

problematic for shared work practitioners. They frequently rely on their networks of contacts to help them locate high-quality experts and documentary knowledge sources, and they rely on their knowledge of who “authored” knowledge contributions to assess quality. Expertise-seeking novices have potentially great difficulties at all stages of knowledge reuse. They are often unable to define the search question properly, they have difficulty locating and judging the quality of knowledge sources, and they may lack the ability to apply expert answers and advice. In the worst case, they may not even know that they need expertise or that relevant expertise exists. Secondary knowledge miners are often completely divorced from the sources of the knowledge they try to reuse. However, they have analytic expertise that most novices lack. If they use a disciplined methodology, their chances of successful knowledge reuse may be greater than those of the typical novice.

These observations provide a basis for explaining and predicting the success of IT support in different knowledge reuse situations and for making normative recommendations about how to improve reuse success. The next section examines the creation of knowledge repositories as one important factor in the success or failure of knowledge reuse.

The Role of Repositories in Knowledge Reuse

THE PREVIOUS SECTION HAS SHOWN that different types of knowledge reusers need different things from knowledge repositories. Shared work producers and shared work practitioners need contextualized knowledge (although the former can more effectively use raw records than the latter, who require more “sanitized” and quality-checked records). Expertise-seeking novices needed decontextualized knowledge, knowledge about what contextual information is useful, and help recontextualizing the information for their unique settings. Secondary knowledge miners can benefit from, but usually have to make do without, in-depth knowledge of the contextual influences on the creation of records stored in repositories.

However, the different types of knowledge reusers do not always get what they need from repositories, for reasons that have to do in part with how repositories are created. In this section, I argue that the contents of repositories tend to differ when knowledge producers knowingly create records for different types of users. (In some cases, these differences may help the consumers. In other cases, they may not.) In the subsequent section, I argue that a great deal of effort is required to produce repositories that meet users’ needs, and that knowledge producers, who are frequently expected to produce high quality repositories, often lack both the motivation and the resources to do so. Therefore, successful knowledge reuse requires providing proper incentives for the knowledge producers and shifting some of the burden of packaging and disseminating knowledge onto intermediaries.

The purpose and content of records in repositories often differ depending on whether the record keepers are knowingly documenting *only for themselves*, for *others who are similar to themselves* in work product or practice community, or for *others who*

are dissimilar to themselves in knowledge and outlook, such as members of a different department in the same organization, novices, or customers.

Documenting for Ourselves

Most knowledge workers create records for their own personal use [36, 41, 45], to remind themselves of details they will need later. For example, in a study of the organization of knowledge workers' offices and desks, Malone [45] found that both files and "piles" were used as organizing devices, and that one important function of piles was to remind people of work that needs to be done.

Because they are intended to serve as reminders, personal records reflect what the knowledge worker expects will be useful later. Since knowledge workers often have difficulty anticipating distant future needs for information [29], their records tend to be biased toward short-term needs. For example, Axline [9] found that ERP implementation teams keep records of "issues" that need resolution, because they will have to resolve such issues before the project is done. However, they tend not to keep records of things with longer-term value (such as details of why issues were resolved as they were, which are useful when a new team member is brought on board, or when the system is upgraded). For such longer-term matters, they rely on their own memories. Unfortunately, they often forget, and the organization loses access to knowledge when team members depart [9].

Depending on note taking skill and diligence, the records knowledge producers do keep may be voluminous, as Orlikowski [61] found of some of the software support workers she studied: "I find it helpful for myself to put in as much information as possible. I find that the more explicit I was earlier, the more it helps me remember when I go back to work on the incident" [61, p. 8].

But other knowledge workers keep sketchier notes and find themselves having difficulty using them later, as was the case of the chairman of an invention proposal review committee (called a Technology Assessment Panel) at Xerox PARC, who was often not able to review his notes until a month or more after the review meeting, at which time he found his notes difficult to interpret [57].

Records kept by knowledge workers for their own use are often quite informal. In her study of software support workers, Orlikowski [60] observed that there is a difference between the kinds of notes one makes for oneself and the kinds one makes for public consumption: "When we used personal notes before [instead of a publicly accessible database] I wouldn't have to worry about [making sure that my entries are technically accurate and correct] because I knew nobody else had to look at that" [60, p. 10].

The exception to this general rule is records that are to be retained for a long time or for legal reasons, in which case the records are more likely to be more formal and sanitized: "[I]f the group is required to store the information for a long period because the process takes a long time (as with medical conditions) or because of legal requirements (as with doctors), the memory will be much more formalized" [1]. In

addition, personal records often contain details and contextual knowledge that are stripped away when the records are prepared for use by others [1].

In short, people documenting for themselves tend to produce documentation as a by-product of the work itself. The documentation they produce is of two types. One is a raw record of notes and communications, often informal, containing shorthand detail that makes sense only to the author (when she or he can remember what it means). This record is biased toward what is likely to be useful to the author in the short term. The second is finished or in process working documents, which may, but often do not, contain much information about the rationale for various decisions.

Documenting for Similar Others

When people knowingly create documents that they know others will read, they consciously or unconsciously shape their records into public documents. When the others for whom they are writing are quite similar to them in terms of knowledge, the extent of the shaping can be relatively minor, because the readers can be expected to be familiar with much of the general and specific knowledge that went into producing the record. And, having similar goals and interests, they can generally be trusted to reuse the information in acceptable ways. Therefore, less effort is required to shape the public “face” of the document.

For example, physicians writing entries in medical records can keep their records quite sparse, because they rely on the knowledge they share with colleagues, and they build on the previous entries in a patient’s chart:

[Physicians] are sensitive to the inferences that can be drawn from particular items. They can rely upon those inferences not only to include information which might otherwise seem relatively trivial, but to exclude particular items (or even categories of object) knowing that any competent reader would be able to make sense of the entry and retrieve the relevant information.

[Entries are] not so much a précis of what happened in the consultation, but rather a brief sketch which provides a certain sense of the event. . . . The very brevity of the entry, the omission of certain categories of item, coupled with the presence of some mentioned treatment, serve as an embedded instruction to the reader to turn to previous entries in order to retrieve the relevant information. The practices that doctors use to assemble the records are the selfsame practices on which they rely in reading the records. [35, pp. 356–358]

On the other hand, the technical support personnel at Zeta [61] were anxious to elaborate the details of a support encounter so that their peers could quickly understand the situation and not have to create their own extensive documentation: We are trying to document so other people can benefit. . . . If you do document well then typically the next person doesn’t have to document again [61, pp. 8, 10]. At the same time, they censored their records in a way they would not have to do if the records were only for them. They

were very aware that the process documentation was publicly available . . . at least to everyone in [the department] and possibly to other departments and offices in the future. Reflecting on this awareness, they monitored and censored what they did and did not enter into incident histories. . . . I am always concerned about being politically correct, professional, diplomatic. . . . [I]t's very easy to put in some sarcastic comments about a person, . . . but I've always not done that specifically for the reason that a year or even six months from now that person may see that incident and take offense and it could jeopardize future relations. [61, p. 10]

This type of “sanitizing” is not too difficult, and it is generally the case that documenting facts and procedures for similar others is not too taxing. However, documenting *rationales* even for knowledgeable colleagues can be very challenging indeed. Gruber and Russell [29] explain:

[R]ationale explanations cover a broad range of information requests. . . . [T]he design rationale author has to anticipate the space of questions that might be asked by the reader, and formulate possible rationale explanations in advance [of the questions being asked]. . . . [Details] captured at the time of design will answer only a fraction of the questions asked by designers about existing designs. [29, p. 340]

Furthermore, even when documenting facts and procedures, there is a limit on how much knowledge can be documented:

I'll use myself as an example. . . . I've been doing research on knowledge and learning for five years. I can write a document that can give you some ideas, but realistically there is no way I can write something to make you understand the subject the way I do. At some point you need to bring in the expert. [69]

Because documenting for similar others eventually becomes burdensome and of diminishing value, many knowledge management systems provide access both to expertise and to experts. For example, Booz Allen's Knowledge On-Line (KOL) system provides access to the detailed resumes of every employee's experience and areas of expertise in addition to documents about consulting engagements [69]. And Ernst & Young's Knowledge Web provides consultants with access to a variety of different kinds of documents, including unfiltered work products (edited primarily to add keywords and maintain client confidentiality) and highly filtered “Power Packs” of proposals and deliverables about particular topics [39].

Documenting for Dissimilar Others

When people knowingly document knowledge for others who are very dissimilar to themselves—such as people in other departments, novices in an area where the documenters are experts, and external customers—two issues come into play. The first is awareness that the reusers lack not only general or technical knowledge but

also the ability to understand the relevance (and irrelevance) of specific or contextual knowledge. The second is awareness that reusers may misuse explicit knowledge.

Documenting for dissimilar others often involves removing from explicit documentation the detailed contextual knowledge that the expert requires but that the novice does not know how to use. For example, Ackerman found that documenting knowledge for novices often meant “re-authoring” the available information:

[A]uthors of the database often remarked that they needed to rewrite the questions and answers.

This was true for a number of reasons. The most commonly cited reason was to make the questions and answers more generalized. The lack of generalizability resulted from several causes. In their questions, users might include many details, only some of which were relevant, because they did not understand the problem at hand. This was particularly true for novices in a subject domain. The answerer, on the other hand, might have used implicit knowledge in his response. For example, the answerer might have known that the asker had only a SUN workstation in his office or that the asker knew how to run “chkdsk” on his hard disk.

Additionally, the database author might know that the answer could be made more general in order to answer more questions. This might involve abstracting both the question and the answer. Occasionally, the author would feel it necessary to correct incorrect, incomplete, or incoherent answers.

The removal of contextual information, including the writer’s implicit knowledge of the reader, is required to make the information understandable across organizational boundaries. [1, p. 7]

In addition, people who document for dissimilar others must think about how those others could use or misuse the information. For example, fashion designers may go to extraordinary lengths to develop precise body measurements to inform their own design work [63]. But, when communicating with suppliers, they find that it is possible to introduce errors by giving suppliers too many measurements: In attempting to satisfy every requirement, “they ‘lose control over the silhouette’ of the garment” [63, p. 224]. Solutions to this problem involve providing less specification, specifications of what the designers did *not* want, a modified physical garment, or drawings that suggest critical lines or shapes rather than giving many measurements. Ironically, providing photorealistic detail sometimes hides significant characteristics: “[T]he idiosyncrasy of novel design detail . . . [is] mostly better conveyed by a sketch with the deliberate noting of [details like seams and darts] and even their exaggeration” [63, p. 335].

In other cases, documenting for dissimilar others means taking special pains to ensure that records sound objective and professional. For example, writing about system administrators who worked on an outsourced, contractual basis for a U.S. manufacturing firm, Schultze wrote:

[In an environment where contract workers were vulnerable to replacement and where blaming and fingerpointing were pervasive], the system administra-

tors relied on their documentation to protect themselves against attack. . . . For the documentation to serve as a dependable shield, . . . it had to meet certain standards of complete and accurate reporting. The system administrators believed that they adhered to higher documentation standards than other contractors working in US Company. . . . When documenting events . . . , the system administrators were careful to delineate the role that they and others played in completing a task. . . . [T]hey strove to reconstruct the work scene in as accurate and complete a way as possible, relying on rhetorical devices such as writing in a passive voice and writing about themselves as subjects in their own texts. They also used strikeouts and added text that qualified or corrected earlier text rather than editing it or deleting it. All of these strategies helped create an impression of reliable, objective reporting. . . . Developing such accounts was hard work and both [system administrators] resented it. [66]

Documenting for dissimilar others can sometimes mean withholding information or providing inaccurate information. For example, Grudin notes that “social, political, and motivational concerns [may] prevent the explicit statement of the real reasons underlying design choices” [31, p. 457]. Similarly, Hertzum notes that one of the roles served by documents in professionals’ work is “to share information with some, yet withhold it from others” [36]. He observes:

Access to information may affect the distribution of power and privileges in an organisation. As a precaution against unintended sharing many documents are intentionally ambiguous and thereby understandable to competent readers only. [36]

For example, Orlikowski [61] found that managers and technical support specialists at Zeta were concerned about “inappropriate assignment of blame and use of information out of context.” Therefore, they restricted outsiders’ access to their database:

In response, CSD managers . . . [allowed] restricted access to individuals on the basis of their personal trustworthiness [and they] offered alternative means for obtaining ITSS data, which did not require direct access to the Notes database [such as a weekly extract from the database for a particular set of clients]. [61, pp. 25–26]

Similarly, Malhotra and colleagues found that members of a virtual new product design team “began to realize that they were reluctant to post entries in the repository because they had the impression that entries should be complete before posting; in the words of one team member, ‘this repository might be subpoenaed in the future if there is an accident on the launch pad’” [44]. Although they soon began to focus on sharing knowledge within the team, they restricted access to their project record—originally set up to allow managers as well as team members access—after (the team members felt) management responded inappropriately upon reading some early entries [42]. In general, Ackerman has concluded, “participants . . . [are] much more willing to provide information if that information was not going to be shared beyond the work group” [1, p. 7]. (See Kovalainen et al. [40] for a similar observation.)

Anonymity of records may reduce unwillingness to share information with dissimilar outsiders. For example, workers in a Finnish paper mill commented on the “risk” of publicly documenting events in a paper diary located in the managers’ office. Consequently, the paper diary was “seriously underutilized” [8]. By contrast, an electronic diary accessible from the shop floor was much more extensively used [40]. Researchers credited two factors: first, the diary was viewed as an informal repository that did not replace formal management reports and, second, few of the entries were signed:

An important observation on the comparative paper/electronic entries is that people’s names often appear in the paper diary [and they can be observed going into the foreman’s box to write or read entries], while in the electronic diary they are very rare—unaddressed notes are the norm. [40, p. 52]

In short, when people knowingly create records for the use of others who are quite dissimilar, the records they create will be quite different from the ones they create for themselves. The records may be re-authored to make them more general and to remove contextual information. Accuracy may be sacrificed to promote understandability. Careful attention may be given to the “tone” of records, so that they “sound” official, objective, and fair. In cases where trust is lacking, potentially damaging information may be withheld from documents or false information supplied. Ackerman distilled these observations into the propositions:

The shorter the distance the information might travel or the less likely it was that the information could be viewed by strangers, the more informal the information content was likely to be. . . . [Conversely, presenting a publicly acceptable “face” in records] was very important when the content was likely to be viewed by decision-makers or management. . . . As one interviewee joked, “What you say may come back to haunt you.” [1, p. 7]

When people have purposely created records for themselves, they may strenuously resist making the records public.³ Authors of records known to be publicly available may require assurances of anonymity before they can be induced to contribute. Note, however, that anonymity of records reduces the ability of other users to judge the quality of contributions, a need identified in an earlier section. Thus, it seems likely that it is not possible to satisfy the needs of all potential users with a single repository. In keeping with these observations, one of the design guidelines offered by Dixon [22] for the situation of “serial reuse” (analogous to my “shared work producer” situation) is that the learning record of the team *not* be shared with or forwarded to others.

Summary

Because the different types of knowledge reusers (shared work producers, shared work practitioners, expertise-seeking novices, and secondary knowledge miners) need different things from their knowledge repositories, the quality and contents of their

knowledge repositories are important factors in the success of knowledge reuse. Therefore, it is important to identify the factors that influence the quality and contents of knowledge repositories. Two major factors in the quality and contents of knowledge repositories are *who authors the entries* and *for whom they author the entries*. The usual expectation is that knowledge producers will author repository documents for reuse by others (whether community-of-practice members or novices). But this expectation contains two problems. First, the records knowledge producers make purposely for their own use are not likely to meet the needs of others. Second, the records knowledge producers make for others may not meet their own needs, and therefore, they may not have adequate incentives to produce quality documents that meet the needs of others.

These observations imply that people who are interested in ensuring successful knowledge reuse need to play close attention (1) to the *costs* involved in creating good repositories, (2) the *incentives* knowledge producers have to contribute to repositories for use by others, and (3) the need for, and roles of, *human and technical intermediaries* in the “repurposing” of repositories developed by knowledge producers to make them appropriate for use by others (and in facilitating other aspects of reuse). In the next section, I address these three issues.

Toward Successful Knowledge Reuse

AS DIXON [22] POINTS OUT, successful knowledge transfer or reuse requires a complete solution. It is not just a matter of providing access to information technology and repositories. It also means careful attention to the design of incentives for contributing to and using repositories and to the roles of intermediaries to develop and maintain repositories and to facilitate the process of reuse. Incentives and intermediaries are important, owing to the costs of making good repositories and using them.

The Costs of Making and Using Good Repository Records

The challenges of inducing people to document their work have been noted in many walks of life. For example, sales reps often delay filling out and handing in customer contact reports or computer programmers often shirk program documentation. In addition, finding relevant explicit knowledge is inherently problematic, even when one has documented and stored the explicit knowledge oneself. We have all had trouble retrieving particular documents from our files and deciphering our own notes after a lapse of time. The difficulties stem, in part, from the incompleteness of records and the problems of indexing. Indexing represents a serious challenge to the reuse of knowledge, even by the knowledge documenters themselves [12, 36].

It has been observed that people are most likely to produce documentation when the documentation is intended to benefit themselves rather than others, when the benefit is immediate rather than delayed, and when the effort required is minimal, as when the documentation is produced as a by-product of the work itself [30, 31]. But even when

these conditions are met, the effort required to produce and use good documentation can be prohibitive. Majchrzak and colleagues [44] provide a vivid example. They studied a virtual team involved in new product design. The team used a collaboration technology with sophisticated keyword search and retrieval facilities. At the outset of the project, team members jointly agreed to document everything that might be of use to them later and to assign at least three keywords to each document in the database. The team immediately became overwhelmed by the demands of documentation and soon settled for less documentation and more synchronous interaction:

The team quickly discovered that there was too much information being generated to be captured and that much of the information was likely to have only transient utility . . . ; thus, their expectations that they would document everything was just too cumbersome. As a result, the project team lead began to call for twice weekly brainstorming sessions using teleconferencing coupled with the Internet Notebook. . . . In preparation for each meeting, team members would post not complete entries detailing the work they had done but instead would post incomplete entries which would then be the source of much discussion during the teleconference. [44, pp. 10–11]

Further, key word-indexing discipline broke down almost immediately:

A core capability of the tool the team was provided was the ability to reference-link entries and apply multiple keywords, and then use powerful search capabilities to identify similar entries. . . . However, . . . the team never found it desirable to use reference linking, multiple keywords, and more than rudimentary search capabilities. . . . For example, only 37% of the entries had two or more keywords. [44, pp. 12–13]

Thus, even in situations where knowledge producers are making records for their own reuse, the costs of producing a high quality repository may be too high. When the knowledge producers are building repositories for use by others, the costs of creating documents and indexing them for reuse escalate. Consider these examples of the shared work practitioner situation.

At Zeta, Orlikowski found “unanticipated problems”:

with the amount of time available to produce high quality and sanitized knowledge for dissemination. . . . It’s not even just writing it. . . . I’m on the review committee and that’s where a lot of time is as I’ve got to review every document. . . . The delay in implementing this second knowledge dissemination plan was due to . . . the lack of resources to provide high quality, sanitized knowledge for consumption by customers. [61, p. 24–25]

Similarly, at AMS, time constraints were seen as an inhibitor to quality contributions:

We have half achieved success with AMSCAT (AMS’s Center for Advanced Technologies) . . .

If you ask people, they will tell you that they really want to learn and they really want to contribute, but they are out working on a project for 15, 16, 17 hours a day, five to six days a week, and knowledge management is not their first priority. [6, p. 12]

At Andersen Consulting, the task of providing adequate “context” for reusable documents in the Knowledge Exchange was felt to be burdensome:

Contextualizing documents would involve getting both professionals and knowledge management staff to write brief summaries for each document, describing the context in which the knowledge was generated and used, where it may be useful, and where it should—and should not—be used. This would certainly add an extra burden to the many professionals contributing documents, perhaps reducing the likelihood that they would bother to contribute their knowledge. Moreover, it would require a major effort to contextualize the existing hundreds of thousands of documents. [38, p. 10]

Indexing repository entries is also a problem in the shared work practitioner situation. For example, knowledge managers at Brøderbund found that they had to redo their document classification scheme because of lack of consistency in indexing:

We had to overhaul the [knowledge repository] after three years because nobody was following a consistent style [of classifying documents]. . . . [P]eople were building their case bases with different parameter settings, so it became like a soup of knowledge, and nobody could find anything. [25]

McKinsey & Company also experienced challenges with the indexing of documents in their consulting databases [47]. They found that all documents could be accurately classified and retrieved on the dimension of “industry” (of the client for which the engagement was done), but classification by “function” was problematic: “[W]hat one consultant calls ‘organizational design,’ another may call ‘change management.’ . . . How does one capture and classify all the consulting experience in the firm when people refer to the same things in different ways?” [47, p. 27] McKinsey eventually solved the problem by encouraging a self-organizing and evolutionary classification process:

[Instead of creating a limited list of keywords that most people could buy into], we split one team into multiple teams, each one responsible for a “sub-domain” of knowledge. We encouraged each sub-team to form communities . . . [and] to devise its own map of terms for the knowledge. . . . What started as an electronic library evolved into a dynamic and living array of both codified and people-based knowledge. [47, p. 29]

McKinsey’s elegant solution would not fully address the problems in the novices-seeking expertise situation, however, since novices often lack knowledge of appropriate key word search terms. Thus, even more effort is required to make a repository easily searchable by novices.

Incentives to Make and Use Good Repository Records

The costs involved in creating and using repositories would not necessarily be a problem if they were balanced by appropriate incentives. But in many cases they are not. Grudin [30] made this point generally for systems that support collaborative work, including shared work product and shared work practice situations. Ackerman [1] particularizes Grudin's observation for the organizational memory systems that are the focus of this paper. Ackerman says organizational memory systems

are subject to the issue of incentives. Grudin . . . has pointed out the complexities of upstream versus downstream costs in adopting and using group and organizational systems. To the extent that OMS require upstream costs, such as those in indexing, and to the extent that the downstream payoffs are unclear, they will likely fail. The use of organizational memory adds to the cost. Not only is there the cost of storage and indexing, there may be additional costs in retrieval and interpretation of the information. [1, p. 5]

The situation in which people contribute to a knowledge repository that both they and others can use has been described as the "discretionary database problem" [49]. The discretionary database problem has frequently been observed in the shared work practitioner situation such as the use of knowledge repositories implemented in consulting firms. The challenge is to ensure that people make high quality contributions without free riding on the efforts of others, particularly when people are pressed for time or competing with each other on the basis of performance. Orlikowski [60], for example, found that consultants in the hectic, up-or-out environment of the firm named Alpha did not use Lotus Notes as extensively as expected for knowledge sharing. Lack of appropriate incentives to contribute was implicated in the failure [60].

Conversely, providing appropriate incentives has often been proposed as a solution to the discretionary database problem. In many consulting firms, explicit attention to extrinsic rewards has been found necessary to promote contributions to shared knowledge repositories. At Booz-Allen: "Consultants contribute for many reasons. But two reasons stand out: the system enhances their work and enhances their reputation among their colleagues" [46, p. 5]. But both of these motivations are reinforced by the explicit reward systems in the firm:

The job of persuasion is made easier by the firm's reward structure. Some practices offer awards for excellent contributions. These range from tombstones for the consultant's desk to the CMT practice's award of \$50,000 to a winning team for them to spend on learning-related activities of their own choice. The only restriction is that they must report back their learning to the practice as a whole. But more significantly, developing intellectual capital has become one of the four criteria used when determining promotion and bonuses. [26, pp. 8-9]

Nevertheless, however useful they are, explicit rewards may be insufficient in the face of unsupportive organizational norms:

Although all Ernst & Young consultants are aware of and evaluated on their use of and contribution to the knowledge management system, there were still some concerns about whether Ernst & Young truly has the type of open, sharing culture which is essential to the optimal operation of its knowledge management system. John believed the current resistance may be partially based on contributing individuals' over-concerns about client confidentiality or else mistrust that their information will be used correctly. On the receivers' end, there may also be some resistance to relying on the professional judgment of a "stranger." [39, p. 18]

In short, one way to overcome the disincentives created by the high costs of creating and using knowledge repositories, particularly those designed for use by others, is to provide appropriate organizational incentives and back them up with appropriate organizational norms. But two observations can be made about this solution. First, it is likely to be less successful the greater the knowledge distance between the intended user of the knowledge and the knowledge producers. Producers have the greatest natural incentives to create repositories that benefit themselves directly in use. They have some but lower natural incentives to create repositories for similar others (in the shared work practice situation), where they can potentially benefit from others' reciprocity. They have lowest natural incentives to document for dissimilar others, where the primary reward is the user's gratitude.

Second, appropriate incentives and norms may not be sufficient to overcome disincentives to create and use repositories, because the costs are so high. In her study of two ERP implementation teams (an instance of the shared work producers situation), Axline [9] found that both teams tried repeatedly to create repositories but generally failed to use them successfully later. Their failures could be attributed in part to absence of incentives such as pressure from leaders and clients. But part of the failure could be attributed to technological inaccessibility, absence of human support for technology use, and lack of knowledge about when, why, and how to document successfully. Overcoming these challenges would have required skilled intervention. Put differently, successful knowledge reuse, even in the "self use" case, may require the use of intermediaries.

The Role of Intermediaries in Knowledge Documenting and Reuse

It takes more than information technology to document and reuse knowledge. It takes "organizational work" [13]:

[T]he knowledge produced [by a community of practice] doesn't readily turn into something with exchange value or use value elsewhere. It takes organizational work to develop local knowledge for broader use. [13, p. 99]

And much of that organizational work is currently done most effectively by human intermediaries. An important question for future research is the degree to which the roles of intermediaries can be undertaken by technology. In the meantime, it is important to

understand the roles that intermediaries can, should, and do play in successful knowledge reuse.

The role of various types of intermediaries, facilitators, and knowledge engineers has been studied off and on in the information systems (IS) field since Culnan's [17] pioneering work on information "chauffeurs." IS research has examined the role of intermediaries and facilitators in the context of electronic meetings systems [28, 54] and of electronic communication systems [59, 62]. As one sifts through these and other empirical studies, it becomes clear that facilitators and intermediaries play a still underappreciated role in knowledge reuse.

Some of the activities that can be performed by intermediaries to facilitate knowledge reuse have already been described in the examples presented earlier in this paper: abstracting, indexing, authoring, sanitizing, and so on. A few additional examples will serve to illustrate the breadth and magnitude of intermediaries' roles.⁴

In a widely cited knowledge management project, Hoffman-LaRoche successfully streamlined its new drug application process to the FDA. The application documents, which range up to 200,000 pages each, take many months to prepare, which delays the receipt of revenues from newly developed drugs [67].

Using carefully developed criteria for document evaluation . . . outside consultants found that Roche's new drug applications did not always communicate key messages and sometimes included contradictory, ambiguous, and inappropriate information. [67]

Armed with this realization, a knowledge management project team was able to develop a knowledge map that now helps employees see how their knowledge fits in and what kinds of knowledge are needed in a successful drug application. The result was a dramatic shortening in "time to market" for new drugs.

Content in knowledge repositories must often be filtered to ensure high quality, and it must be pruned when it becomes out of date. These are tasks that knowledge producers often dislike. Effective performance often requires intermediaries. The case of Booz-Allen is particularly instructive in this regard:

Accompanying the new technology were new practices to ensure that the knowledge generated in client projects was captured and shared. Twelve knowledge managers, dedicated to particular practices, work with teams to ensure that, where appropriate, sanitized and generic lessons learned are entered into the system. There is a strong emphasis on quality control. As one partner says: "You don't want just anything to get posted there. We definitely have a different view at Booz-Allen than some others who suggest you let everyone post anything and you create a free market. Absolutely not! That wastes everyone's time: all knowledge is not created equal."

[T]he norm, overall, is still to use KOL to deliver "know-what" rather than "know who." One downside of this is that besides often being misleadingly incomplete, documents can become outdated on KOL. In response, BAH is making a concerted effort to periodically prune the system. [26, pp. 8, 12]

British Petroleum introduced video-teleconferencing and other tools to promote the formation of “virtual teams”—communities of practice that would help each other with such tasks as tricky oil field maintenance jobs. A key aspect of their knowledge management intervention was the role of “coach”—when the coach was absent, people often did not understand how and why they were expected to use the new technology:

A subgroup of the core team . . . was responsible for helping participants understand both how to use [the Virtual Teamwork] technology and how it could further their work. This effort was deliberately called “coaching” rather than “training.” . . . The core team was . . . convinced that extensive coaching was essential to the success of the project . . . An unplanned event helped prove them right. Due to budget constraints, one of the projects . . . was set up without coaching. . . . This project was the only one of five that failed. The problem [was that the team members lacked] an understanding of why they should bother [using the technology]. [37]

It has often been observed that latent communities of practice do not always become active or self-sustaining. “[T]hey require significant investments of time and effort . . . and often need a dedicated facilitator” [65]. For example, at Buckman Labs, “systems operators” played important roles in ensuring that people made high quality contributions to the K’netix system. The “sysops” monitored discussions, tracked down answers, provided for language translations, gave positive feedback to contributors, and published codes of ethics and appropriate behavior:

System operators (Sysops) were appointed to monitor the discussions in the forums, track requests and make sure they were answered. Sysops would try to get answers in 24 hours; if not they would contact people directly and ask them to respond. Additionally, they were to give positive feedback to those who did respond. Since there were likely to be cultural differences and sensitivities, Sysops were to monitor the content of messages. . . . Three translators were hired and Sysops would decide which messages were to be translated into English with technical replies to be translated back to the originator’s own language. The goal for completion of translation was 48 hours. [14, p. 8]

Another role played by facilitators in knowledge reuse is content authoring. One technique that a number of companies have found useful in their knowledge management efforts is the learning history, a written document that records a successful team’s “thinking, experimentation and arguments in a way that forced them to reflect on their experiences” [65, p. 43]. These 20-page documents, “intended primarily for the organization as a whole” [65, p. 44], go “far beyond a mere list of ‘best practices’ or ‘process improvements’” [65, p. 43]. Not surprisingly, the assistance of intermediaries is required to produce these “knowledge objects”:

A small group of internal staff members and outsider learning historians “distill” the raw material into a coherent set of themes and write the report. We have chosen the word “distill” carefully to convey the essence of this activity—taking

volumes of data from interviews and then purifying and refining the “raw data” into a form that the organization can hear. [65, p. 51]

When organizations appreciate the value of facilitators and intermediaries in successful knowledge reuse, the level of effort applied can be very great. At Ernst & Young for example, “there are close to 200 professionals working at the CBK [Center for Business Knowledge], and it offers a call center, supplementary business research and analysis functions, coordination of the knowledge networks, and administration for the firm’s databases of internal and external information” [39, p. 10].

Three additional observations should be made about the role of intermediaries in successful knowledge reuse. First, even with strong organizational support, intermediaries may experience demanding jobs, lack of cooperation, and lack of appreciation. This was the case at McKinsey in the early days of that company’s knowledge management efforts:

Making the new practice coordinator’s position effective proved more challenging. Initially, these roles were seen as little more than glorified librarians. It took several years before the roles were filled by individuals (often ex-consultants) who were sufficiently respected that they could not only act as consultants to those seeking information about their areas of expertise, but also were able to impose the discipline necessary to maintain and build the practice’s data bases. [53, p. 5]

Second, even with the commitment of many organizational resources, some tasks likely to facilitate effective knowledge reuse may remain undone. For example, at Booz-Allen, the knowledge repository was believed lacking in the synthesis of knowledge from many different projects:

It was widely felt, for example, that what KOL was missing were the lessons learned not from individual projects—the system was excellent at capturing those—but across multiple projects. “The primary limitation of the system at the moment is the fact that it is based on project output. You’ve got lots of reasonably disparate pieces of information about different projects, but if you were interested in a topic, say, data warehousing, no one has synthesized all the pieces of work we’ve done in that area and put it down in one place.” [26, p. 13]

Third, intermediaries can play important roles not only in the “discretionary database” situations where knowledge is documented for others to use, but also in situations where people are creating documentation for their own later reuse. In the case of self-reuse, as discussed earlier, the incentives for knowledge documentation are believed to be strongest. But even here the incentives may not be strong enough to overcome the costs, particularly in the face of time pressure. Consequently, intermediaries are recommended to promote knowledge reuse in the case of shared work product teams:

The [aerospace new product design] team never did resolve the issues of speedy knowledge retrieval. In the end, they recommended that [virtual] teams should

consider establishing a role of a Knowledge Manager. Such a Knowledge Manager can serve several functions. First, the knowledge manager can ensure that valuable information is not left unrecorded in the knowledge repository by reviewing the roadmap of the repository and identifying obvious gaps in logic. For example, if explanations, circumstances, and constraints for quantitative estimates are missing from an entry (for example, the circumstances under which the impinging holes in a design would be too expensive), the Knowledge Manager could ask for more detail. Second, the knowledge manager can help to ensure that the entries in the information repository can be reviewed by outsiders, by providing an easy way for others to get the information they need. Finally, the knowledge manager can ensure that the team is able to make use of the documentation that they create by reminding the team of past information and helping them find it when needed. While we believe that such a position has merits, we believe there is significant research needed to determine how to best organize knowledge for re-use. [44, pp. 15–16]

Summary

Earlier sections of this paper have argued that different types of knowledge reusers have different needs from their knowledge repositories and that the quality and contents of knowledge repositories depend in part on who creates the entries and for whom they are creating the entries. In this section, I have examined three additional factors in the successful reuse of knowledge in repositories: the costs involved in creating and using entries, the incentives people have to create and use entries, and the roles of intermediaries in the creation and maintenance of repositories and the facilitation of their use. The conclusion is that intermediaries can play an important role in successful knowledge reuse even when the incentives to create and use knowledge repositories are greatest (when the knowledge reuser was the knowledge producer). The future challenge for the IS field is to find ways for information technology to take on an increasing share of the intermediary role.

Conclusion

THIS PAPER HAS BEEN A FIRST ATTEMPT to pull together evidence from a wide variety of sources into the rudiments of a theory of successful knowledge reuse. What might at first glance appear to be a unified phenomenon is actually quite varied. Even without factoring in the roles of intermediaries, there are at least four distinct knowledge reuse situations involving different types of knowledge reusers: shared work producers, shared work practitioners, expertise-seeking novices, and secondary knowledge miners. Each type of knowledge reuser has different needs from repositories and can be expected to encounter different kinds of problems when attempting to reuse knowledge. Consequently, successful knowledge reuse is in part a matter of designing repositories that meet reusers' needs. But, it is also a matter of deciding who should

author (or re-author) the repositories, how to provide incentives for high quality contributions, and what role should be assigned to human and technical intermediaries for managing repositories and facilitating their use.

One plausible conclusion from this analysis is that repositories created by one group for one purpose are unlikely to be successfully reused by other groups for different purposes without considerable rework or other kinds of intervention. How much and what kinds of intervention are required may depend on the "knowledge distance" between the knowledge reusers and the original knowledge producers. To what extent intervention can be accomplished by technical (as opposed to human) intermediaries is an important avenue for future IS research.

Acknowledgments: This paper has benefited greatly from the comments of Steve Frenkel, Ann Majchrzak, Mark Silver, Varun Grover, and an anonymous reviewer.

NOTES

1. Dixon and I concur that successful knowledge transfer involves a combination of interventions such as a business driver, use of information technology, a knowledge sharing process, appropriate incentives, process facilitation, and so on.
2. Consultants, of course, may also be members of a team working on the same project, for example, the shared work producer situation described above.
3. In the expert systems literature, it is often assumed that any resistance of experts to knowledge engineering has to do with fears of job or power losses. This paper suggests an alternative explanation for such resistance, based on the possibilities that public access to knowledge bases affords for misinterpretation, misuse, and reprisals against the authors.
4. Consult Dixon (2000) for even more examples.

REFERENCES

1. Ackerman, M.S. Definitional and contextual issues in organizational and group memories, available from the author's web site, www.ics.uci.edu/~ackerman, 1994.
2. Ackerman, M.S., and Halverson, C. Considering an organization's memory. *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work*. New York: Association for Computing Machinery, 1998, pp. 39–48.
3. Ahuja, M.K., and Carley, K.M. Network structure in virtual organizations. *Journal of Computer-Mediated Communication*, 3, 4 (1998), available at www.ascusc.org/jcmc/vol3/issue4/ahuja.
4. Alavi, M. Managing organizational knowledge. In R.W. Zmud (ed.), *Framing the Domains of IT Management: Projecting the Future Through the Past*. Cincinnati, OH: Pinnaflex Educational Resources, 2000, 15–25.
5. Alavi, M., and Leidner, D.E. Knowledge management systems: issues, challenges, and benefits. *Communications of the AIS*, 1, 7 (February, 1999), available at cais.aisnet.org.
6. American Management Systems: The knowledge centers. Case No. 9-697-068. Boston: Harvard Business School, 1998.
7. Anand, V.; Manz, C.C.; and Glick, W.H. An organizational memory approach to information management. *Academy of Management Review*, 23, 4 (October 1998), 796–809.
8. Auramaki, E.; Robinson, M.; Aaltonen, A.; Kovalainen, M.; Liinamaa, A.; and Tunna-Vaiska, T. Paperwork at 78kph. *Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work*. New York: Association for Computing Machinery, 1996, pp. 370–379.

Appendix: An Example of Reuse Requiring Contextualized Information

Five years ago we were involved in a major proposal made by an international information and communications systems company to a US aerospace manufacturer. . . . We chose computer conferencing as the process to integrate the proposal activities, which lasted four months. We designed the conferencing and collaboration environment (the electronic, virtual workspace) to . . . Enable cross-functional collaboration amongst the ten specialist skill groups comprising the proposal team [and] Co-ordinate a team activity which was geographically distributed. . . .

[A]t the height of the proposal development . . . over 90 percent of the team's communication was taking place within computer conferences [and] Everyone was using the overall Project Management computer. . . . Material from all conferences was being used to create the emerging proposal document.

The team generating the knowledge was the team that used the knowledge in their work, . . . That knowledge was immediately reviewed and applied within the particular skill group and was available for impact assessment across the whole team. Discussions about how to proceed at critical points was team-wide, and final decisions were recorded for all to read. . . .

[Later] Senior management created a task force (mostly of people not previously involved) to take the final proposal document and genericize it into a proposal template for aerospace bids—similar to the process of documenting methodology discussed earlier.

This was not a success. The resulting document was more of a testament to the efficiency of the Replace All word processor function than an item of usable intellectual capital. The result of the proposal work was still there, but it was not on its own accessible. It contained little context for the proposal (which in fact had a two year history); its structure was set for the original customer and could not be made generic without losing its logic; and nothing in it gave any clue as to why the team had chosen any of the solutions they proposed.

Soon after, however, we ourselves became involved in an engagement with another aerospace customer. . . . One of the first things we did was to open the computer conferencing archive that had been created by the earlier proposal team. The following happened:

The new team read the conferences. They understood not only the details of the proposal the original team had created, but why they had done it, and what went on while they were doing it. . . .

- They were able to abstract key points from the captured process of the earlier proposal. These were both points that they already knew they needed answers to, and also points they didn't know they needed answers to.
- They were able to contact people whose expertise was apparent from the recorded conferences, and consult with them about the new proposal.

- They were able to place the genericized aerospace proposal template in context, and were able skillfully to pick and modify items which were relevant to their proposal.

Source: J. Gundry and G. Metes. *Team Knowledge Management: A Computer-Mediated Approach*. Malmesbury, UK: Working by Wire Working Paper, 1996, www.knowab.co.uk/wbwteam.html.

9. Axline, S. Proactive Adaptation in ERP Implementation Teams: Mechanisms of Team Learning. Ph.D. dissertation, Claremont Graduate University, 2000.
10. Bashein, B.J., and Markus, M.L. *Data Warehouses: More Than Just Mining*. Morristown, NJ: Financial Executives Research Foundation, Inc., 2000.
11. Bashein, B.J.; Markus, M.L.; and Finley, J.B. *Safety Nets: Secrets of Effective Information Technology Controls*. Morristown, NJ: Financial Executives Research Foundation, Inc., 1997.
12. Blair, D.C. The management of information: basic distinctions. *Sloan Management Review*, 26, 1 (Fall 1984), 13–23.
13. Brown, J.S., and Duguid, P. Organizing knowledge. *California Management Review*, 40, 3 (Spring 1998), 90–111.
14. Buckman Laboratories (A) and (B). Case No. 9-800-160, 9-800-033. Boston: Harvard Business School, 1999.
15. Choudhury, V., and Sabherwal, R. Information technology and the boundary of the firm: a knowledge-based perspective. In R.W. Zmud, A. Segars, and J. Sampler (eds.), *Redefining the Organizational Roles of Information Technology in the Information Age: Workshop Proceedings*. Cincinnati, OH: Pinnaflex, 2001 forthcoming.
16. Conklin, E.J., and Burgess-Yakemovic, K. A process-oriented approach to design rationale. In T.P. Moran, and J.M. Carroll (eds.), *Design Rationale: Concepts, Techniques, and Use*. Mahwah, NJ: Lawrence Erlbaum, 1996, pp. 393–427.
17. Culnan, M.J. Chauffeured versus end user access to commercial databases: the effects of task and individual differences. *MIS Quarterly*, 7, 1 (March 1983), 55–65.
18. Davenport, T.H. If only HP knew what HP knows. Ernst & Young Center for Business Innovation, Cambridge, MA, n.d., available at www.businessinnovation.ey.com.
19. Davenport, T.H.; De Long, D.W.; and Beers, M.C. Successful knowledge management projects. *Sloan Management Review*, 39, 2, (1998), 43–57.
20. Davenport, T.H.; Jarvenpaa, S.L.; and Beers, M.C. Improving knowledge work processes. *Sloan Management Review*, 37, 4 (Summer 1996), 53–65.
21. Dean, D.L.; Lee, J.D.; Pendergast, M.O.; Hickey, A.M.; and Nunamaker, J.F. Enabling the effective involvement of multiple users: methods and tools for collaborative software engineering. *Journal of Management Information Systems*, 14, 3 (Winter 1997–1998), 179–222.
22. Dixon, N.M. *Common Knowledge: How Companies Thrive by Sharing What They Know*. Boston: Harvard Business School Press, 2000.
23. Duffy, D. It takes an e-village. *CIO Enterprise Magazine*, October 15, 1999.
24. El Sawy, O.A., and Bowles, G. Redesigning the customer support process for the electronic economy: insights from storage dimensions. *MIS Quarterly*, 21, 4 (December 1997), 457–483.
25. Elliott, S. Brøderbund builds strong “case” for internal, external knowledge sharing. *Knowledge Management in Practice*, Fourth Quarter. Houston, TX: American Productivity and Quality Center, 1998, available at www.apqc.org.
26. Galunic, C., and Weeks, J. Managing knowledge at Booz-Allen & Hamilton: knowledge on-line and off. INSEAD, Fontainebleau, France, May 2, 1999.
27. Gottschalk, P. Knowledge management systems: a comparison of law firms and consulting firms. *Informing Science*, 3, 3 (2000), 117–124.
28. Griffith, T.L.; Fuller, M.A.; and Northcraft, G.B. Facilitator influence in group support systems: intended and unintended effects. *Information Systems Research*, 9, 1 (March 1998), 20–36.
29. Gruber, T.R., and Russell, D.M. Generative design rationale: beyond the record and replay paradigm. In T.P. Moran and J.M. Carroll (eds.), *Design Rationale: Concepts, Techniques, and Use*. Mahwah, NJ: Lawrence Erlbaum, 1996, pp. 323–349.
30. Grudin, J. Why CSCW applications fail: problems in the design and evaluation of organizational interfaces. *Proceedings of the Conference on Computer-Supported Cooperative Work*. New York: Association for Computing Machinery, 1988, pp. 85–93.
31. Grudin, J. Evaluating opportunities for design capture. In T.P. Moran, and J.M. Carroll (eds.), *Design Rationale: Concepts, Techniques & Use*. Mahwah, NJ: Lawrence Erlbaum, 1996, pp. 453–470.
32. Gundry, J., and Metes, G. Team knowledge management: a computer-mediated approach. *Working by Wire SM TM*. December 1996, available at www.knowab.co.uk/wbwtteam.

33. Hackbarth, G., and Grover, V. The knowledge repository: organizational memory information systems. *Information Systems Management*, 16, 3 (Summer 1999), 21–30.
34. Hansen, M.T.; Nohria, N.; and Tierney, T. What's your strategy for managing knowledge? *Harvard Business Review*, 77, 2 (March–April 1999), 106–116.
35. Heath, C., and Luff, P. Documents and professional practice: “Bad” organisational reasons for “good” clinical records. *Proceedings of Computer Supported Cooperative Work '96*. New York: Association for Computing Machinery, 1996, pp. 354–363.
36. Hertzum, M. Six roles of documents in professionals' work. *ECSCW'99: Proceedings of the Sixth European Conference on Computer Supported Cooperative Work*. New York: Association for Computing Machinery, 1999, pp. 41–60.
37. Knowing the drill: virtual teamwork at BP. *Ernst & Young Center for Business Innovation*, n.d., www.businessinnovation.ey.com...al/issue1/features/known/body.html, accessed February 8, 1999.
38. Knowledge management at Andersen Consulting. Case No. 9-499-032. Boston: Harvard Business School, 1999.
39. Knowledge management at Ernst & Young. Case M-291. Graduate School of Business, Stanford University, 1997.
40. Kovalainen, M.; Robinson, M.; and Aurmaki, E. Diaries at work. *Proceedings of the Conference on Computer Supported Cooperative Work*. New York: Association for Computing Machinery, 1998, pp. 49–58.
41. Lansdale, M. The psychology of personal information management. *Applied Ergonomics*, 19, 1 (1988), 55–66.
42. Majchrzak, A.; Rice, R.E.; Malhotra, A.; King, N.; and Ba, S. Technology adaptation: the case of a computer-supported inter-organizational virtual team. *MIS Quarterly*, 24, 4 (2000), 569–600.
43. Majchrzak, A.; Rice, R.E.; Malhotra, A.; Nelson, K.; and Ba, S. Knowledge-sharing in a creative inter-organizational virtual team. Working paper, University of Southern California, December 1999.
44. Malhotra, A.; Majchrzak, A.; Carman, R.; and Lott, V. Radical innovation without collocation: a case study at Boeing Rocketdyne. *MIS Quarterly* (2001 forthcoming).
45. Malone, T.W. How do people organize their desks? Implications for the design of office information systems. *ACM Transactions on Office Information Systems*, 1, 1 (January 1983), 99–112.
46. Manasco, B., and Perelman, L. Booz Allen's global knowledge strategy. *Knowledge Inc.: The Executive Report on Knowledge, Technology & Performance*, 1, 1 (n.d.), available at www.knowledgeinc.com.
47. Manville, B. A complex adaptive approach to KM: reflections on the case of McKinsey & Company, Inc. *Knowledge Management Review*, 8 (May–June 1999), 26–31.
48. Markus, M.L. Electronic mail as the medium of managerial choice. *Organization Science*, 5, 4 (November 1994), 502–527.
49. Markus, M.L., and Connolly, T. Why CSCW applications fail: problems in the adoption of interdependent work tools. *Proceedings of the Conference on Computer-Supported Cooperative Work*. New York: Association for Computing Machinery, 1990, pp. 371–380.
50. Markus, M.L.; Majchrzak, A.; and Gasser, L. A design theory for systems that support emergent knowledge processes. In R.W. Zmud, A. Segars, and J. Sampler (eds.), *Redefining the Organizational Roles of Information Technology in the Information Age: Workshop Proceedings*. Cincinnati, OH: Pinnaflex, 2001 forthcoming.
51. McDermott, R. Why information technology inspired but cannot deliver knowledge management. *California Management Review*, 41, 4 (Summer 1999), 103–117.
52. McDonald, D.W., and Ackerman, M.S. Just talk to me: a field study of expertise location. *Proceedings of the ACM 1998 Conference on Computer-Supported Cooperative Work*. New York: Association for Computing Machinery, 1998, pp. 315–324.
53. McKinsey & Company: Managing knowledge and learning. Case No. 9-396-357. Boston: Harvard Business School, 2000.
54. Miranda, S.M., and Bostrom, R.P. Meeting facilitation: process versus content interventions. *Journal of Management Information Systems*, 15, 4 (Spring 1999), 89–114.

55. Moorman, C., and Miner, A.S. Organizational improvisation and organizational memory. *Academy of Management Review*, 23, 4 (October 1998), 698–723.
56. Moran, T.P., and Carroll, J.M. (eds.). *Design Rationale: Concepts, Techniques, and Use*. Mahwah, NJ: Lawrence Erlbaum, 1996.
57. Moran, T.P.; Chiu, P.; Harrison, S.; Kurtenbach, G.; Minneman, S.; and Melle, W.V. Evolutionary engagement in an ongoing collaborative work process: a case study. *Proceedings of the ACM 1996 Conference on Computer-Supported Cooperative Work*. New York: Association for Computing Machinery, 1996, pp. 150–159.
58. O'Dell, C., and Grayson, C.J. If only we knew what we know: identification and transfer of internal best practices. *California Management Review*, 40, 3 (Spring 1998), 154–174.
59. Okamura, K.; Fujimoto, M.; Orlikowski, W.J.; and Yates, J. Helping CSCW applications succeed: the role of mediators in the context of use. *Information Society*, 11, 3 (July–September 1995), 157–172.
60. Orlikowski, W.J. Learning from Notes: organizational issues in groupware development. *The Information Society*, 9, 3 (July–September 1993), 237–250.
61. Orlikowski, W.J. Evolving with Notes: organizational change around groupware technology. CISR WP No. 279, Sloan WP No. 3823, CCS WP No. 186. Center for Information System Research, Sloan School of Management, MIT, Cambridge, MA, June 1995.
62. Orlikowski, W.J.; Yates, J.; and Okamura, K. Shaping electronic communication: the metastructuring of technology in the context of use. *Organization Science*, 6, 4 (July/August 1995), 423–444.
63. Pycocock, J., and Bowers, J. Getting others to get it right: an ethnography of design work in the fashion industry. *Proceedings of the ACM 1996 conference on Computer Supported Cooperative Work*. New York: Association for Computing Machinery, 1996, pp. 219–228.
64. Ricks, T.E. Lessons learned: army devises system to decide what does, and does not, work. *The Wall Street Journal*, May 23, 1997, A1, 10.
65. Roth, G., and Kleiner, A. Developing organizational memory through learning histories. *Organizational Dynamics*, 27, 2 (Autumn 1998), 43–59.
66. Schultze, U. A confessional account of an ethnography about knowledge work. *MIS Quarterly*, 24, 1 (1999), 43–79.
67. Seemann, P. Knowledge management at Hoffmann-LaRoche. Ernst & Young Center for Business Innovation, Cambridge, MA, n.d., available at www.businessinnovation.ey.com.
68. Selling knowledge on the “Net.” Ernst and Young Center for Business Innovation, Cambridge, MA, n.d., available at www.businessinnovation.ey.com.
69. Tristram, C. Common knowledge. *CIO Web Business*, September 1, 1998, webbusiness.cio.com, accessed June 23, 2000.
70. Walls, J.G.; Widmeyer, G.R.; and El Sawy, O.A. Building an information system design theory for vigilant EIS. *Information Systems Research*, 3, 1 (1992), 36–59.
71. Wenger, E. *Communities of Practice: Learning, Meaning, and Identity*. Cambridge: Cambridge University Press, 1998.
72. Wenger, E.C., and Snyder, W.M. Communities of practice: the organizational frontier. *Harvard Business Review*, 78, 1 (January–February 1999), 139–145.
73. Yates, J.; Orlikowski, W.J.; and Okamura, K. Constituting genre repertoires: deliberate and emergent patterns of electronic media use. *Academy of Management (Best Papers Proceedings 1995)*, 353–357.
74. Zack, M.H. Managing codified knowledge. *Sloan Management Review*, 40, 4 (Summer 1999), 45–58.