PERSONALIZING KNOWLEDGE DELIVERY SERVICES: A CONCEPTUAL FRAMEWORK

ANN MAJCHRZAK and RAMNATH K. CHELLAPPA
Department of Information and Operations Management
Marshall School of Business
University of Southern California
Los Angeles, CA 90089-1421
majchrza@usc.edu and ram@marshall.usc.edu

LYNNE P. COOPER
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109-8099
Lynne.P.Cooper@jpl.nasa.gov

ALEXANDER HARS
School of Law, Business and Economics
University of Bayreuth, Geb. B9
95440 Bayreuth, Germany
alexander.hars@uni-bayreuth.de

The research described in this paper was carried out under contract with the Jet Propulsion Laboratory, California Institute of Technology and the National Aeronautics and Space Administration.

November 5, 2002

PERSONALIZING KNOWLEDGE DELIVERY SERVICES: A CONCEPTUAL FRAMEWORK

ABSTRACT

Consistent with the call of the Minnesota Symposium for new theory in knowledge management, we offer a new conceptualization of Knowledge Management Systems (KMS) as a portfolio of personalized knowledge delivery services. Borrowing from research on online consumer behavior, we describe the challenges imposed by personalized knowledge delivery services, and suggest design parameters that can help to overcome these challenges. We develop our design constructs through a set of hypotheses and discuss the research implications of our new conceptualization. Finally, we describe practical implications suggested by our conceptualization – practical suggestions that we hope to gain some experience with as part of an ongoing action research project at our partner organization.

PERSONALIZING KNOWLEDGE DELIVERY SERVICES: A CONCEPTUAL FRAMEWORK

INTRODUCTION

Alavi & Leidner (2001) review the literature discussing applications of IT to organizational knowledge management initiatives and identify three common applications: 1) the coding and sharing of best practices, 2) the creation of corporate knowledge directories, and 3) the creation of knowledge networks. They argue "IT as applied to KM need not be constrained to [these] certain types…because the advances in communication and information technologies enable greater possibilities than existed with previous classes of IS" (p. 115). We use this quote to ground our focus: discussion of a new class of KMS.

KMS are focused in part on encouraging knowledge reuse (Alavi & Leidner, 2001), where reuse is defined as the application of others' knowledge to a particular domain or problem in a way that improves the quality or efficiency of the reusers' work activities (Majchrzak, Cooper & Neece, in press). Thus, when the right knowledge is reused at the right time, benefits such as decreased coordination costs, less rework, fewer repeated mistakes, less time to invent, etc. are achieved.

We are focused on reuse of others' knowledge to solve non-routine problems, e.g., problem-solving that requires flexibility, ingenuity, innovation, and creativity. Examples of business processes that involve such non-routine problem-solving include organizational design, strategy development, innovation, new product development, and troubleshooting problems that rarely repeat because they arise from interactions among large numbers of highly complex and interdependent components (such as rocket launches). Markus et al. (2002) refer to these processes as EKPs (emergent knowledge processes) because they involve a swirl of actors, actions, processes, and knowledge that emerge in chaotic ways. Reuse with EKPs often involves the creative application to a new context of an idea initially used in another context. For example, Majchrzak et al. (in press) describe how a scientist was able to reuse newly acquired knowledge about how the British textile industry reduces static cling generated by its fabrics to develop an

innovation in a scientific instrument designed to assess electrostatic properties on Mars. Because EKPs are fundamentally different than routine problem-solving processes, the process by which knowledge is reused in an emergent work process is different than the process by which knowledge is reused in a routine problem-solving process, a point we have demonstrated empirically elsewhere (Majchrzak et al., in press; Faniel & Majchrzak, 2002).

Given that knowledge reuse in EKPs is different than knowledge reuse for routine problem solving, it is perhaps not surprising that KMSs supporting knowledge reuse for EKPs might need to be different than KMSs supporting knowledge reuse for routine problem-solving. The most frequently mentioned distinction in the KM literature is between KMS that support networks of individuals versus KMS that support knowledge repositories (such as lessons-learned or best practice or corporate repositories), with the latter focused on helping EKPs and the former focused on helping more routine problem-solving (Hansen et al., 1999; Brown & Duguid, 1991). For example, Hansen et al. (1999) describe knowledge repositories as valuable for a consulting project in which knowledge objects are being reused (such as repetitive consulting), while network support is valuable for connecting consultants to each other as they develop business strategy for clients.

According to Markus et al. (2002), this distinction is too simplistic. Simply providing networks to support EKPs is unlikely to be sufficient, in part because the innovator needs not only a network that connects her to others in her community of practice, but also emergent access to new ideas from a range of knowledge repositories and communities. Therefore, research on KMSs to support EKPs must go beyond simply supporting communities of practices. However, there is a relative dearth of research attention on such KMSs (Grant, 1996). In addition, according to a survey by Davenport et al. (1998) of KMS initiatives in industry, there are far fewer industry initiatives as yet on KMSs to support innovative work and thus we have few industry exemplars to guide theory-building and research inquiry.

In this paper, we present a framework for conceptualizing KMSs that support the reuse of others' knowledge in non-routine problem-solving – a support that is designed to provide the reuser with the right knowledge at the right time, and, as a consequence, help the EKP reuser to generate better solutions in less time. In the spirit of the Minnesota

symposium, this framework is still under development, evolving as part of an action research project with a partner organization. We believe this framework to be of value to the broader community of KMS researchers since it suggests new research questions. This paper proceeds by first presenting an overview of our conceptualization, and why we choose the conceptualization we did, followed by a more detailed description of the conceptualization and the implications for further KM research and practice.

OVERVIEW OF CONCEPTUALIZATION: PERSONALIZATION OF KNOWLEDGE DELIVERY SERVICES

EKPs provide unique challenges for the design of KMS. First, EKP work requires synthesizing across a variety of different and often unpredictable inputs; thus, EKP workers need KMS support not just to identify and transfer a single pre-specifiable knowledge object, but also to be exposed to a variety of inputs and to synthesize across these inputs. The communities of practice approach to KMS (Brown & Duguid, 1991; 2001) does not provide a conceptual understanding of how to connect people across communities. Moreover, the expert directories approach assumes the reuser can define the expert or the area of expertise needed. Yet, our empirical research (Faniel & Majchrzak, 2002) as well as conceptual research by Markus (2001) indicates that the nature of EKP work often means that the reuser may not be able to identify the experts who can shed light on a problem; moreover, the experts, when found, may have difficulty matching their expertise with the problem.

A second challenge imposed by EKPs is that the reuser cannot be assigned the entire responsibility for reuse. The reuser will not be able to always pre-define when she needs reuse, what she needs, where to find it, how to evaluate what she has found, and how to reuse it if she decides to adopt it. For example, in the process of defining the problem itself or in searching for reusable ideas, the reuser may redefine where to search, search criteria, and new problem definitions (Majchrzak et al., in press). Moreover, reusers have very little time. Faniel & Majchrzak (2002) and Majchrzak et al. (in press) have found that, when faced with time limitations, EKP workers will simply not reuse others' knowledge and will make their decisions without the benefit of others' input. Thus, expecting reusers to pre-define their reuse needs such as by proactively searching

knowledgebases, crafting questions for discussion boards or communities of practices, or searching expert directories is an expectation ill-suited to the way that EKP work is done.

A final challenge of EKP work for KMSs is the critical importance of contextualization. Our previous research on EKP workers indicate that they need to know a range of quite detailed information (such as parameters, assumptions, constraints) about the context of the source knowledge (knowledge they are intending to use) (Cooper, in press; Majchrzak & Faniel, 2002). However, for contextual information about the source to be of value in determining if the source knowledge should be reused, EKP reusers reinterpret or reconstruct the source's contextual information within their current context. For example, if the current context is characterized by an interaction of two components, the reuser will examine the source context for evidence of those two components (or analogous components) even if the two components weren't central to the source's previous problem context (Faniel & Majchrzak, 2002). This reconstruction is often iterative, as the reusers may not initially realize that the interaction of the two components is critical to their problem solution, or that the source has any knowledge of the two components and their interaction. A KMS that supports this reconstruction process is needed.

These challenges suggested to us to conceptualize KMSs not as a single system that needs to be built, but rather as a set of "knowledge delivery services" that configure, package, and deliver knowledge to EKP workers as the need arises. Services are defined as activities and mechanisms by which knowledge is identified, interpreted, manipulated, and then reused. Services may include a community of practice discussion forum from which a member gets advice (Orr, 1996) or a knowledge repository with a query mechanism for finding information about previous lessons-learned reports, or a facilitator for a design meeting to help the team identify and match patterns between a team's evolving design rationale and previous or other design activities.

The challenges imposed on reuse by EKPs also suggested to us that knowledge delivery services need to be personalized to the individual user, available knowledge sources, the nature of the problem, problem-solving stage, and the work context such as the nature of the organization and team dynamics. In concept, then, such personalization can foster the reconstruction and recontextualization of a source's knowledge for the

unique problems the EKP workers face. Such personalization can also expose the EKP workers to multiple sources of knowledge as they need it and suggest ways to synthesize. Finally, such personalization can require less effort by the reuser for the entire reuse process by suggesting, guiding, and prompting EKP workers about reuse opportunities.

There is limited guidance in the KMS literature for what a Personalized Knowledge Delivery Service (PKDS) would look like. However, by examining findings on the personalized delivery of online products and services from consumer behavior research, we can begin to address this gap. In the consumer behavior research, personalization has meant tailoring a service to "suit the customer's needs" (Solomon, et al., 1985). The ability to dynamically tailor such services online for one customer relies on techniques in which the understanding of a particular customer's preference is predicted in advance based on an understanding of other similar customers and use contexts (Mobasher, et al., 2000). Hence one can argue that using a similar procedure that can predict in advance the knowledge needs of one reuser based on information about similar others and contexts will lead to greater reuse. Thus personalization technologies lend themselves to encouraging reuse.

It is therefore important to understand how the technologies behind personalization lend themselves to the goal of reuse. Extensive developments in personalization for online consumers have resulted in two well-defined techniques: recommendations based on decision rule systems, and those based on collaborative and content based filtering systems (Mobasher, et al., 2000). Of particular interest to knowledge services is the operation of the filtering systems. These systems not only utilize a customer's own information, but they also incorporate preferences and profiles of others through a correlation engine, and more recently these systems are also incorporating pattern discovery tools into their portfolio. Such techniques are vital in personalization of knowledge services, e.g., an alert or pointer to a particular knowledge source could be sent to a manager in one meeting based upon patterns developed from another manager who has a similar user profile working within a similar problem profile. Thus knowledge reuse is not simply restricted to a particular knowledge source or even a problem.

CONCEPTUAL FRAMEWORK

Before we seek to understand mechanisms through which knowledge delivery can be personalized, an understanding of common personalization techniques in the consumer context is warranted. Personalization in the physical world has always been a part of a merchant-consumer relationship, and research in consumer behavior and retailing refers to personalization as the manner in which "service employees relate to customers as people" (Mittal and Lassar, 1996). Prior to the emergence of electronic commerce, personalization was limited to in-store interactions between consumers and salespeople. Today, firms are able to collect information on not only formal transactions but also non-transaction behavior such as browsing and sampling. This new information combined with other pre-existing information on customer preferences has allowed firms to create rich user profiles leading to personalization of online products and services (Cingil, et al., 2000). These user profiles are essentially data that are interpreted to constitute the interests of a user (Chellappa and Sin, 2002; Kramer, et al., 2000; Cingil, et al., 2000), where the data may include anything from subjects of interest to the user, the number of times a resource is used, and the cumulative time a user has spent using a website or resource. Online firms use this profile information to personalize products and services to fit customer's stated needs as well as perceived needs. Borrowing from the consumer literature, then, we define personalization of knowledge delivery services as:

> *The use of technology and user information (including context) to tailor the delivery of knowledge from a knowledge source (which may be an individual, an organization, a repository, or a set of repositories) to the user, so as to increase decision making effectiveness*

This definition suggests that a knowledge delivery service can be considered to be personalized if the knowledge and/or its delivery is altered to fit the user's stated as well as inferred needs based on the user information that has been either previously obtained or provided in real-time. Our definition suggests that just as online personalization is dependent upon a merchant's ability to accurately create customer profiles (Mulvenna, et al., 2000), personalization of knowledge services is based on the construction of an accurate knowledge profile of the knowledge source, the reuser, and the problem or task

at hand. In the customer context, personalization technologies categorize objects (typically Web pages) and subjects (consumers) and perform matches between and across objects and subjects and determine a set of actions to be recommended for personalization (Mobasher, et al., 2000). Similarly, we propose that a personalized knowledge service is delivered as determined by appropriate matching of four elements: user characteristics, knowledge source characteristics, work-context characteristics, and problem characteristics:

1. Characteristics of user: A knowledge user profile is the information about the attributes and preferences of the individual and may include an individual's role on a design team, time pressure, commitment to solving the problem, problem-solving style, and current knowledge domains.

2. Characteristics of knowledge sources: A source may be a knowledge object (such as an idea) or a group of knowledge objects (such as a document or database). A knowledge source profile is the information about the attributes of various knowledge sources and include a) ideas and artifacts such as documents, hyperlinks to other sources, stories, or names of individuals, b) assessments of the quality, reliability, potential obsolescence, etc about the ideas and artifacts, and c) metadata (such as date, keywords) associated with each source.

3. Characteristics of work context: A work context profile is the information about the attributes of the work context experienced at the time that the knowledge is being delivered and may include time pressure, whether the knowledge is being used as part of a team effort, stage in project lifecycle, stage in team's problem-solving capability, team cohesiveness and openness of knowledge-sharing, amount of knowledge-overlap among team members, and the incentive structure encouraging risk-taking and innovation (versus deliberate and conservative decision making).

4. Characteristics of the problem: A problem profile is the information abut the attributes of the problem at the time that the knowledge is being delivered and may include how well-defined the problem is at the time, whether the problem-solving process is at a point when it will benefit more from convergent or divergent thinking, and the uniqueness of the problem.

These characteristics are not constant but change as work and time progress. Ideally the knowledge delivery system monitors such changes. As it updates the profiles of the four elements, a knowledge delivery system searches for opportunities to inject knowledge. Such opportunities may be, for example, a team member making reference to how a solution used in a previous project would be helpful for the current problem.

This raises an opportunity for a personalized knowledge delivery service to inject knowledge about that previous project. The conditions and the way in which the knowledge is to be delivered is depicted in the "dynamic delivery profile". Thus, for example, a project manager walking through her day may have many different dynamic delivery profiles constructed for her by the knowledge delivery service as she shifts from an exploratory meeting with colleagues about a new proposal, to a status meeting about an existing project, to completing some administrative paperwork at the end of the day. In the exploratory proposal meeting, her dynamic delivery profile might identify new ideas based on the words proposal developers use in developing their ideas, while during the review meeting, the profile may identify convergent types of knowledge (e.g., keeping track of action items), while at the end of the day when she is filling out administrative paperwork, the profile may include partially filled out administrative templates based on the way she has completed the administrative forms in the past.

When a knowledge injection opportunity occurs (such as a team member making reference to how a solution used in a previous project might be helpful for the current problem), the dynamic delivery profile can be brought to bear to construct a personalized knowledge delivery service. For example, the knowledge about a previous project may be delivered as a voice-synthesized statement by the computer about the previous project and an analysis made by the computer about how the previous project relates to the discussion at hand. Or, the knowledge may be delivered as an alert flashed on the desktop (or laptop) of the team member indicating a hyperlink to the files for that referenced project. Table 1 presents some examples of personalized knowledge delivery services.

Table 1: Examples of How Knowledge Delivery Services (in italics) can be Personalized for a Context
▪ Provide *alerts* about the latest developments in a user's discipline ▪ Provide a *decision map* to meeting facilitator describing the options being considered ▪ Provide project manager at end of meeting with *meeting notes linked* to outside knowledge sources with pertinent information ▪ If in divergent phase, inject new knowledge sources, ask "what if questions", and suggest alternative problem definitions through such services as *interactive queries*, or *decision rationale methodologies*. ▪ If in convergent phase, provide *templates* & knowledge sources such as *boilerplates* to help with preparing proposal, refining ideas, preparing for reviews

- Provide a *table* depicting facts versus stories about a particular event
- Offer *recommendations* to contact certain experts
- Offer *models and simulations* with information on boundary limitations
- Provide *notification* of, or *metrics* about others' reuse of one's own knowledge

Having been offered the knowledge, the reusing team as a whole or individual reusers can choose to use the injected knowledge or not. For example, the team members could verbally inform the delivery service to provide more detail about the past project, request a future reminder/redelivery, or ignore the service's input. How the users react to the knowledge delivery service provides new information that is then incorporated into the dynamic delivery profile. For example, if certain injected knowledge is repeatedly ignored, delivery of such knowledge in the future should either take a different form, or not be delivered at all.

This discussion of this process of a personalized knowledge delivery service suggests that the services must include not only the results (e.g., a voice message or alarm) but also the logic that is necessary to detect a triggering situation. Thus, a PDKS must have two components: one that monitors the environment for triggers and one that generates the outcome. Figure 1 depicts our framework for a Personalized Knowledge Delivery Service.

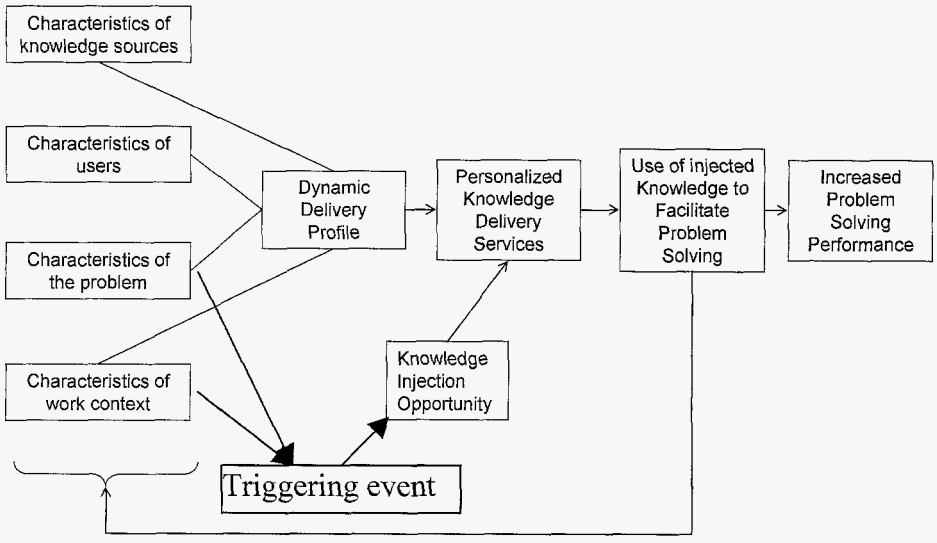

Figure 1: A Framework for Personalized Knowledge Delivery Services

DILEMMAS OF PERSONALIZED KNOWLEDGE DELIVERY SERVICES

On the surface, the practical and theoretical implications of the framework we have briefly outlined above are obvious. The more an organization provides personalized knowledge delivery services for its EKP workers, the more efficiently and effectively the workers will be able to solve their problems. However a review of personalization strategies in the consumer context tells us that employing personalization technologies and deriving the expected benefits are not so simple. Just as with any technological interface, there are obviously issues related to interface design (Kramer, et al., 2000), but the ability to deliver personalization is also based on the ability to acquire and use personal information. Hence the concern of an individual with regards to the collection and use of his/her personal information can significantly affect both the usage and expected benefits from personalization (Volokh, 2000). Therefore an understanding of the online consumer's dilemma, i.e. her tradeoff between value from personalization services and costs from concern for privacy, has been argued to be critical to the success of online personalization strategies (Chellappa and Sin, 2002).

Thus, only if users are willing to share substantial amounts of information *with* the service will the users be able to reap the benefits of substantial amounts of knowledge *from* the service. In the section below, we explore four specific aspects of this information-in vs. information-out tradeoff. By examining these more specific dilemmas, it becomes more apparent how one might handle the information-in vs. information-out tradeoff for personalized knowledge delivery services. The four dilemmas are:

1) The more attributes on which delivery is personalized, the greater the probability that the delivered knowledge will be relevant , but the users may not find enough value given the input burden.
2) The more sensitive the personalized service is to real-time dynamics, the greater the likelihood of providing information when needed, but also when not desired.
3) The more human-facilitated the service, the greater the personalization, but the less ability to identify possible knowledge from others that could be injected.
4) The more information about the user that is obtained to personalize knowledge delivery services, the greater the likelihood that privacy issues will surface, reducing the likelihood of use.

After describing each of these dilemmas in this section, we examine in the following section how each of these dilemmas can be alleviated by suggesting hypotheses for future research.

<u>Dilemma #1: The more attributes on which delivery is personalized, the greater the probability that the delivered knowledge will be relevant , but the workers may not find enough value given the input burden.</u>

On the surface, the feasibility of an idealized personalized knowledge delivery service is questionable for two reasons. First is the amount of information that is needed about the user. The more that a service is personalized; the more information about the user, work context, etc. is required. To obtain this information often involves the user inputting the information. Thus, for a service to be highly personalized to optimize the probability of positively contributing to the decision process, the system runs the risk of requiring so much input from the user that the user will choose not to use the service. Thus, a highly personalized knowledge delivery service may not be perceived to be highly useful, even though it might provide insights about new knowledge.

A second issue with the feasibility of a personalized knowledge delivery service concerns the problem of creating a dynamic delivery profile that matches the four profiles of user characteristics, knowledge sources, problem, and work context. In reuse for routine problem-solving, similarity of task between the source and recipient is an important predictor of reuse (Argote, 1999; Szulanski, 2000). With innovation, however, often matching is done on characteristics of a problem domain that are subtler, or possibly orthogonal to obvious characteristics of a problem. For example, Hargadorn & Sutton (1997) describe IDEO's matching between problem requirements and knowledge sources as based less on the similarity of requirements and more on similarity of concepts. In this way, for example, the inject and eject design of videocassette recorders can be reused for designing portable computer docking stations (Hargadorn & Sutton, 1997) or washing machine dynamics can be reused in the design of rocket engine injector systems (Majchrzak et al., 2000). Thus, the problem of determining the contextual cues that should be the basis on which knowledge injection opportunities are personalized is not entirely clear. So, the dilemma is that: if personalization services are structured to closely match knowledge sources to the knowledge problem, there is likely to be less innovation, providing little additional value than what the individual could have done alone. Thus, a personalized knowledge delivery service might not be perceived to be useful because it either generates non-innovative ideas, or took too much time to identify knowledge that the users were already capable of finding themselves.

Dilemma #2: The more sensitive the personalized service is to real-time dynamics, the greater the likelihood of providing information when it is needed, but also of providing information that is not desired.

Personalization may be based on either real-time information or static information. In consumer settings, static personalization involves using characteristics of the consumer known prior to interacting with the online website. For example, 1-800-FLORIST uses static information provided by the consumer at an earlier point in time about family birthdays to send an email to the consumer reminding her that it might be a good time to consider sending flowers. There is a difference between personalization based on static information and those based on current information. Personalization

abilities may be limited when static rule-based algorithms constructed only on the basis of pre-existing information are employed, as opposed to the deployment of sophisticated pattern development technologies based on current information that can provide significantly more personalized offerings (Mobasher et al., 2000).

A consumer may encounter different flavors of real-time personalization in merchant-consumer settings. One flavor is when real-time information is coupled with static information to provide real-time recommendations. In fact, what often appears to the consumer as "real-time" personalization may rely purely on rule-based algorithms built around static recommendations to the consumer, or collaboratively filtered information obtained from similar other's preferences. For example, Amazon tracks your click stream so that when you consider purchasing a book (even before you put it in the shopping cart), it suggests other books that you might want to consider buying. The second flavor of real-time personalization is truly real time, as in the case of customer services such as those offered by Dell, Earthlink and other technology related firms, in which customer services representatives constantly monitor chat forums and provide real-time interjections for specific problems. While real-time personalization is more accurate in predicting to consumers' preferences, obtaining and organizing real-time information requires more resources.

Borrowing from consumer research, knowledge delivery services may be easier to implement as static than real-time personalization since the personalized delivery can be pre-planned. Example pre-planned personalized services might include user profiles that are used to generate alerts about the latest developments in a particular discipline. However, the problem with static personalization for knowledge reuse in emergent knowledge processes is that the configuration of problems, contexts, and individuals is often so dynamic as to make static-based personalizations less valuable. For example, an alert about new knowledge in a particular discipline may be of limited value if the reuser has already solved her problem. Thus, personalization that responds to real-time dynamic conditions so that new knowledge can be injected as the opportunity arises would intuitively appear to be of greater value than static personalization. For example, a knowledge delivery service that is able to identify the team's stage in the convergent/ divergent problem-solving cycle (Ancona & Caldwell, 1990; Couger, 1996) is more

likely to inject new knowledge at the appropriate time (i.e., when the team is diverging) so that the team finds the knowledge useful than a service that ignores this information.

Although real-time personalization may be more appealing intuitively than static personalization, the dilemma with real-time personalization is that such a real-time delivery may not be suitable in some knowledge delivery scenarios. For example, assuming that the service is able to track each action and conversation, how should the service determine that injected knowledge is welcomed? Knowledge injection may disturb the specific task a user is involved with. An analogy may be drawn with SPAM email that disturbs or causes inconvenience to a customer and has even been equated with trespassing (Quilter, 2002).

Thus, sometimes a real-time personalization provides more value if it is not truly real-time, but is instead, responsive to changing demands at some level of granularity. The dilemma then is how to adjust the sensitivity of the personalization to dynamic conditions of a knowledge process: the less sensitive and static the personalization, the less input required from the user but the greater the likelihood that knowledge is injected too late or inappropriately. In contrast, increasing the sensitivity to dynamic conditions requires more input from the user and may result in a system that "over-reacts", thus creating unnecessary interruptions.

Dilemma #3: The more human-facilitated the service, the greater the personalization, but the less ability to identify possible knowledge from other's that could be injected.

Human-facilitated knowledge delivery services have the advantage of providing highly personalized knowledge. Humans can recognize an event for knowledge injection, attend to the appropriate contextualization cues, search through their memories or databases for appropriate knowledge sources, and translate, configure, and package the knowledge to optimize the likelihood that the reuser will appreciate the benefit of the knowledge. In the EKP domain, there have been many efforts at using human facilitators to encourage reuse. For example: humans have been used to facilitate decision rationale capture and reuse (Moran & Carroll, 1996; Selvin, in press; Conklin & Yakemovich, 1991) and identify knowledge injection opportunities ("opportunity recognizers" by

Leifer et al., 2000 and "technology brokers" by Allen, 1977) for EKPs. The problem with the use of human facilitators is that they can only inject the knowledge they know about or know how to access. Thus, they need to be experts in not only the domain of the problem, but in other domains that might have important input for the problem. Such facilitators are rare and expensive; moreover, if they were asked to facilitate a meeting, experience has found that they quickly prefer to move out of a facilitation role and into a participant role (Cooper, in press). Thus, the dilemma is that, while personalization is greatly enhanced by the use of human facilitators, the greater reliance on such facilitators may decrease the likelihood of identifying new knowledge to inject.

Dilemma #4: The more information about the user that is obtained to personalize knowledge delivery services, the greater the likelihood that privacy issues will surface, reducing the likelihood of use.

For EKP workers, their professional reputations are based on their knowledge (Markus et al., 2002). As a result, they often control what knowledge they convey to others and are often concerned about how that knowledge is used (or misused) by others (Hollingshead et al., 2002; Jarvenpaa & Staples, 2000). In addition, EKP workers are accustomed to significant autonomy in how they do their work; they often chafe at routinization of their core work. For this reason, such well-known structured design procedures such as IBIS have been difficult to diffuse because of the standardization imposed by the language and process by which design decisions are made (Selvin, in press). This need to control one's knowledge and work process suggests that EKP workers are unlikely to welcome being closely monitored while working.

Yet, increasing levels of personalization are only possible with increasing amounts of information about the user. When users are aware that information about them is being obtained (either explicitly or less obtrusively), privacy issues are often raised (Culnan, 2000; Smith, et al., 1996). The more privacy concerns the users have about the information, the less their likelihood of providing the information. For example, for a real-time personalized knowledge delivery service, a facilitator (human or computer-based) would need to "listen-in" on a conversation to know when to inject new knowledge, since the user themselves may not be aware of when new knowledge might be applicable. Having someone or something listen in on one's conversation conjures up

fears of retribution, and insecurity about interpretation. As an example of fear of retribution, the comments of one new product development engineer when talking with his colleagues as they were entering informal entries into a database is telling: "Hey guys, did you know that this database might be subpoenaed one day if there's an accident on the launch pad?" (Malhotra et al., 2000). The implication was, of course, that each entry had a name associated with it, meaning that individual team members could be identified and implicated during the accident investigation if any information in the entry had anything to do with the accident. Retribution may be exercised more subtly. For example, if a knowledge delivery service made a design team aware of all the possible concerns with a design idea – even very low probability ones – and the team made the measured judgment to ignore some of these concerns, can the system be used later to inform a higher authority that the team ignored one of these concerns (especially if that concern later proved to be a cause of product failure)?

In sum, the main concerns that users have are with regards to how the information collected about them for personalization purposes may be used for other purposes. These concerns are related to the perceptions of an individual with respect to their belief regarding a particular activity and how fairly it is conducted (Lind and Tyler, 1988). These concerns about privacy, then, suggest the dilemma that the more personalized the system, the more valuable it might be for the reuser, but the less likely the reuser might be to use the service for privacy fears.

HYPOTHESES FOR OVERCOMING THE 4 DILEMMAS

After making a general case for the adoption of personalized knowledge delivery services, we have raised four dilemmas that need to be resolved in order that personalized knowledge delivery services are successfully deployed. In this section we discuss hypothesized ways of overcoming these dilemmas. Our conceptualization provides 5 hypotheses that can be stated as:

Reuse of others' knowledge to increase the quality and efficiency at which the EKP worker's problems are innovatively solved is more likely when the personalized knowledge delivery services are designed so that:

H1: The services prioritize integration with work over reuse
H2: Contextual information is obtained in a trustworthy environment
H3: Personalization increases over time

H4: Personalization is based on strategic rather than contextual similarity
H5: The services are visibly supported by humans.

Figure 2 depicts a model with these hypotheses.

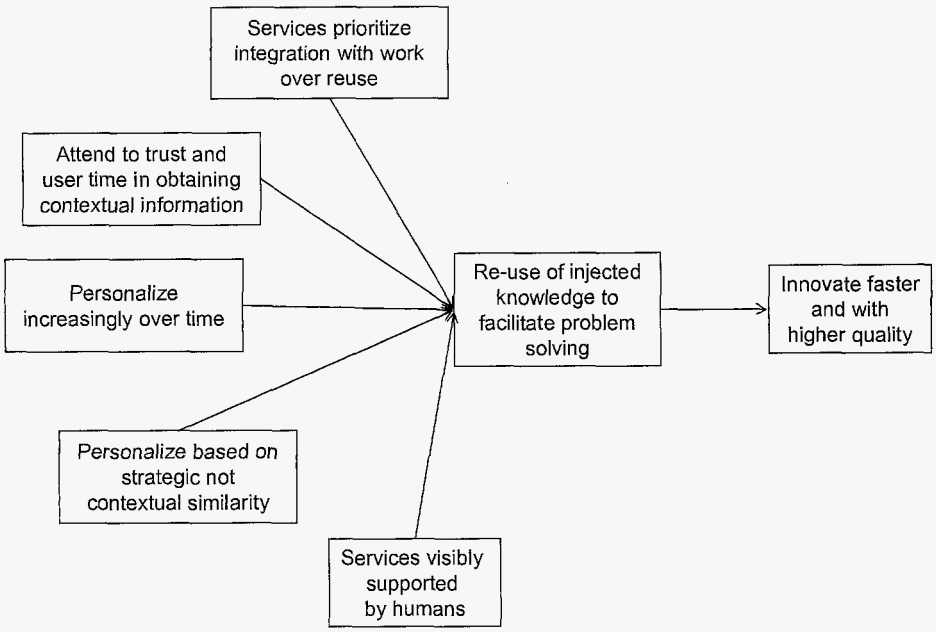

Figure 2: Hypothesized Design Parameters for Effective Personalized Knowledge Delivery Services

In the discussion below, we describe each hypothesized set of design parameters, provide examples of how the design parameters can be implemented within a KMS context, and then describe how we believe our hypothesized design parameters will help to alleviate the dilemmas. Table 2 presents a summary of how each dilemma is helped by the hypothesized design parameters.

Table 2: How each hypothesized design characteristic overcomes different Dilemmas

	Dilemmas

Hypothesized Design Parameters for Personalized Knowledge Delivery Services	#1: Between collecting more attributes to personalize service BUT user not perceiving value given input burden.	#2: Between sensitivity to real-time dynamics allowing more personalization BUT users being interrupted in real-time	#3: Between using humans to facilitate a personalized service BUT humans providing only limited exposure to new source knowledge	#4: Between collecting more knowledge to personalize BUT surfacing privacy issues
#1:Prioritize integration with work over reuse	Bulk of attributes can be pre-defined based on roles, then adjusted for the individual as work is done, not separately, reducing input burden	Services provide advice when advice would be most relevant	Using humans to abstract problems into conceptual [or strategic] queries	NA
#2:Information obtained in a trustworthy environment	NA	NA	NA	Create an environment that alleviates fear of retribution & misuse
#3:Personalization increases over time	Repeated use of personalized knowledge delivery services & feedback reduce need for user input	Repeated use of personalized knowledge delivery services & feedback add value over time by creating more accurate recommendations	NA	NA
#4:Based on strategic rather than contextual similarity	All contextual cues do not need to be captured for valuable knowledge to be injected	NA	Using humans to understand and formulate dimensions of potential strategic similarity	NA
#5:Supported by humans	Review evolving personalization	Review evolving personalization rules to correct	Encourage human facilitators to	Create a panel of individuals who review privacy

	rules to correct macro-level mistakes in interpretation, avoiding the need for each user to make corrections	macro-level mistakes in interpretation, reducing the possibility of undue interruptions	deliver personalized services when supported with computer-based aids to identify new knowledge	and information use policies and serve as the recourse for users who believe their personalized information is misused

H1: A personalized knowledge delivery service is more likely to lead to reuse that improves problem-solving performance when the design prioritizes integration with the EKP, not reuse.

Blackler (1995) and others (e.g., Nidumolu, Subramani & Aldrich, 2001; Schultz & Boland, 2000) have called for the need for knowledge management systems to be integrated into the actual work of the knowledge worker, instead of separate databases or discussion threads that must be searched outside of, or independent of, the work being done. To integrate the KMS into the work suggests that workers should want to use the service to complete their tasks, even if reuse is not the primary goal of a particular task. This suggests then that the goal of the knowledge delivery service cannot be reuse, which is affected by a multitude of factors beyond accessibility of the knowledge; instead it must be innovative EKP performance.

Since the intent of a technology often permeates how a technology is implemented (DeSanctis and Poole, 1994), clarifying the intent of a KMS is important. The distinction between a KMS that supports EKP performance versus a KMS that supports reuse is often not subtle. For example, the intention of lessons-learned databases is reuse, not innovative problem-solving; as such, the databases are designed to require the user to proactively decide when lessons learned are needed, find the database and have defined what lessons to search for. Discussion forums (or communities of practice) are similarly designed to foster reuse through networking not chaotic or radical innovation since they require users to know which community of practice to subscribe to.

Reconceptualizing a lessons-learned database in terms of innovative problem-solving performance suggests that the database needs to be: a) organized in a flexible way since users must flexibly manage their work process, b) tied to existing processes used in

work (e.g., work breakdown structures), c) coupled with a user interface that suggests alternative problem definitions and patterns based on different configurations or elements of the problem (e.g., identifying matches of different configurations of constraints, industry, possible solutions, etc.), d) coupled with an expertise directory (and instant messaging capability or instant discussion forum creation capability) that helps to identify and contact people who can work with the user to further define the problem, and e) coupled with other possible lessons-learned databases so that the search is not necessarily limited to the specific database the user was initially searching. In addition, the system should f) be designed to capture the user's decision making process so that knowledge and ideas of the highest value given the decisions being made are shared. For example, knowledge that negates an assumption made about a design idea is of more value early in the design process than knowledge about more refined cost estimations for implementing the idea. Moreover, when the user has found material that might be useful, the service should g) allow the user to click on the link and be directed not just to text, but to templates, stories, digital artifacts, models, analytic tools or data to manipulate, and h) a community of interested individuals with which to connect to in order to explore these ideas further.

This suggests an integrated set of services for the EKP worker, rather than a set of independent tools: an integration that involves the tools with each other, as well as the tools with the work itself. This also suggests that some services should be provided that lead to improved efficiency of the user's work, even if they don't lead to reuse. For example, providing summarization services such as meeting minutes and action item lists so that the user is able to obtain some demonstrable efficiency gains from the system will succeed at both capturing knowledge for later use in personalizing services and providing the user with efficiency gains immediately.

Services designed with these hypothesized characteristics help to overcome two of the dilemmas. Dilemma #1 is alleviated in part because information for the dynamic delivery profile is being captured not separately from the use context but while work is being accomplished, reducing the burden to the user. Dilemma #2 is alleviated in part because the integration of the tools suggests that the KMS services become natural extensions of the user and her work. Thus, a service that provides the link to an

alternative knowledge source when the user is searching for ideas to solve a problem will welcome the advice rather than view it as an interruption.

H2: Personalization of knowledge delivery services will lead to greater innovative problem solving when user profiles, knowledge source profiles, and problem profiles are constructed from reliable environments and the service is delivered in a trust-worthy environment..

To respond to the privacy concerns raised in Dilemma #4, research in consumer behavior has demonstrated that fear of retribution and misuse of information can be overcome by trustworthy ownership of the environment (Alba, et al., 1997; Schurr and Ozanne, 1985). In particular it has also been shown that the intent to use online personalization systems is dependent upon trust in the online vendor (Chellappa and Sin, 2002). For example a consumer may be less concerned about providing her personal information to a reputed or branded vendor such as Amazon as opposed to a no-name vendor online. In fact the importance of trust has been shown to be especially important in uncertain environments such as Internet based European Community environments (Fung and Lee, 1999).

Trust also reflects a willingness to assume the risks of disclosure (Mayer et al., 1995), and hence consumers who provide personal information during transactions assume the risk of having this information endangered. Such risks are reduced when trust in an entity increases due to activities undertaken by the transacting entity (Koller, 1998). Two important factors that affect a consumer's trust in merchants are their reputation and a consumer's satisfaction with their past interactions. Many researchers in marketing (Doney and Cannon 1997, Gefen 1997) have identified reputation as an important part of trust building when defined as the extent to which consumers view a marketer to be reliable, honest, and trustworthy (Zucker 1986). Empirical evidence also clearly points out that trust follows satisfaction with a service provider (Singh and Sirdeshmukh 2000). In a buyer-seller relationship it has been argued that satisfaction with previous outcomes has a significant impact on trust (Ganesan 1994).

With specific regards to privacy, the Federal Trade Commission provides guidelines on how a vendor should treat customer information. While these guidelines are not legal requirements yet (except in special situations that relate to children and

medical information), the recommendations ask online retailers to notify and provide consumers with the control over how and what information is being collected about them (FTC 2001). The recommendations are in fact closely related to the parameters proposed in theory that have been identified to be contributors to an individual's concern for privacy (Culnan, 2000; Smith, et al., 1996) and related to trust in online environments (Sheehan and Hoy, 2000). Experience with the consumer-side of personalization suggests, then, that a personalized knowledge delivery service must be trustworthy, and that the design elements that are likely to lead to a more trustworthy service is one that has the following parameters: a) notifying users of what information is being collected and how it is collected, b) providing the user with the choice about what information to collect about herself, c) in the case of a misuse of information, provide users with a recourse that corrects such mistakes and avoids future occurrences, and d) informing users about secondary uses of the information. An important element of allaying privacy concerns and creating trust within an organization where information is being collected is factor of control (Culnan 1999, Stone 1983, 1990). Employees within an organization need to possess the belief that they have full control over any monitoring and collection of their transactions. Such beliefs through privacy-allaying practices are important for effective use and success of PKDS initiatives. While people may accept the collection of personal information, they also constantly assess these benefits with the risks of disclosure (Derlega et al. 1993) and hence all practices that reduce the risks of disclosure are critical.

H3: Personalization of knowledge delivery services will lead to greater innovative problem solving when the services are designed to be increasingly personalized over time.

The ability to personalize is based on access to information (Kramer, et al., 2000), and a continuous information acquisition and learning environment can lead to construction of more accurate user profiles and enrich the knowledge source profiles. Thus, during early stages of knowledge services deployment, personalization may be limited, somewhat akin to manual decision rule systems. Such rules are a function of how well the system anticipates current user preferences based on what is known outside

the system (Mobasher et al., 2000). Thus, at best a template may be provided during these early stages. However, with time and feedback, the personalization or correlation engine can correct itself. Such feedback may be explicit such as soliciting the user's satisfaction, or implicit such as observing usage patterns and drawing conclusions. For example, at the beginning, a personalized knowledge delivery service might send alerts on pointers to a database, however if the user does not heed this alert, the system can draw inferences (such as the user not being interested in such alerts) and stop providing that knowledge source.

Repeated use of personalized knowledge delivery services and feedback based on how users behave when receiving the service reduce the burden of user input over time as the service is able to increasingly predict knowledge injection opportunities for the user. This helps in alleviating Dilemma #1. In addition, repeated use and feedback of the services create more accurate recommendations, thus overcoming Dilemma #2's concerns about users being unnecessarily interrupted. Finally, repeated use and feedback also helps in reducing the fears of misuse and privacy discussed in Dilemma #4. Research points out that repeated interactions build trust (Doney and Cannon, 1997) and hence over time an organization has a greater ability to build trust and reduce disclosure related risks to the organizational user of a personalized knowledge delivery service.

H4: Knowledge delivery services designed to personalize based on strategic rather than contextual similarities between source and reuser will lead to greater innovative problem-solving behavior than services based exclusively on contextual similarity.

Contextual similarity refers to the set of contextual cues that have been found in previous research to influence reuse, e.g., similarity between the source knowledge and problem to be solved in terms of goals, organizational context, constraints, problem requirements, as well as similar characteristics between the source and recipient (Szulanski, 2000). Argote's (1999) work on contextual similarity (in which knowledge is viewed as embedded in a network of tasks, actors, organizations, and time and the transfer requires matching on these elements) is an example framework for thinking about transferring knowledge across similar contexts. Using Argote's (1999) framework, then, reuse will be higher when the task, organization, time, and actor are similar between

the best practice to be transferred and the new context receiving the best practice. Personalizing based on such a contextual similarity requires then that these characteristics of the reuse environment be captured and then compared to environments of possible knowledge sources.

As we have argued earlier, limiting comparisons to environments that are similar limits the injection of new knowledge. Thus, we suggest that, instead of looking for contextual similarity, a personalized knowledge delivery service for EKP should look for strategic similarity. By "strategic" similarity, we mean finding low-probability similarities. For example, a cure for Reynaud's syndrome was found by identifying low-probability similarities among different databases (one about fishes, the other about the medical symptoms associated with the disease) (Swanson, 1988; Gordon & Lindsey, 1996). As another example, a new instrument for measuring electrostatic particles on Mars was created by identifying the low-probability similarity between electrostatic measuring devices and development being done in the British textile industry (Majchrzak et al., in press). To find these low-probability similarities require new intelligent text-mining algorithms (e.g., Dworman et al., 2000; Balir & Kimbrough, 2002; Balachandran et al., 1999). These algorithms identify key phrases and concepts used in each group of documents. For example, the phrases used by the members and artifacts for Project #1 are classified as Project #1 concepts and those phrases used in Project #2 as Project #2 concepts. The lists from the two projects are then compared, rank-ordering the concepts by the frequency in which they are observed in both projects. The concepts that are ranked at the bottom of the list as concepts that are *rarely* used in both projects are identified as the low-probability similarities. Some subset of these low-probability similarities constitutes the new knowledge that could be injected into a problem-solving process.

Applying the concept of low-probability similarities to KMS design suggest that a personalized knowledge delivery service may be the continuous running of a low-probability search algorithm (as an intelligent agent) on the evolving database of an EKP work team's deliverables, transcripts of design meetings, discussion forums, and emails evolves over time. Just as a consumer can task a shopbot to keep the consumer informed if price reductions are posted somewhere on the internet for a particular product the

consumer is interested in buying (Clay, et al., 2001), a user or team can task a search algorithm to compare, in the background, their evolving database to hundreds of other evolving databases to identify low-probability similarities. With feedback from the user every time a low-probability similarity is identified, the search engine can become increasingly sophisticated as to the types of low-probability similarities that are of most value to the user.

Designing a KMS to find strategic rather than contextual similarity suggests that all contextual cues do not need to be captured for valuable knowledge to be injected into an innovative problem-solving process . Thus, Dilemma #1 of requiring so many attributes from the user that the value fails to exceed the burden is partially alleviated since not all contextual cues would be needed.

H5: Visible support by humans will facilitate greater use of personalized knowledge delivery services

In the consumer context, it has been argued that human contact is what makes the difference between a "cold and impersonal contact" to a "warm and personal one" (Mittal and Lassar, 1996). In fact, the human interaction element has been always considered crucial during shopping (e.g., "doesn't this suit look great on me?") and hence there is a stress on the quality of interpersonal interaction between the customer and service employee (Bitner, et al., 1994).

In the context of personalized knowledge delivery services, we argue that human support needs to be exercised in three ways. First, human facilitators can be used to deliver the personalized knowledge services such as by keeping track of assumptions that conflict or drawing decision maps of options considered (e.g., Selvin, in press). However, as described in Dilemma #3, humans may not be able to identify new knowledge to bring to bear. In this case, the human facilitator can be supported by the computer-based search algorithms to find low-probability similarities or help identify alternative problem definitions. Additionally, a new role for the facilitator could be defined to expand searching by formulating concept-based queries.

Second, human support is needed in reviewing the evolving personalized knowledge delivery services themselves. Human review of evolving personalization rules may help

to avoid unnecessary time of users correcting service errors, as when the service misinterprets a week of users being concerned about corporate layoffs to users' lack of interest in a service delivery. Thus, human support will help to alleviate Dilemmas #1 and #2.

Finally, human support is critical for ensuring that the services are delivered and maintained in a trustworthy environment (Bitner, et al., 1994), thus helping to alleviate Dilemma #4. For example, a review panel is needed that sets policies and procedures for using information collected about the users in the course of personalizing the knowledge delivery. If users believe these policies are violated, they must be able to turn to the panel to air their concerns (Culnan, 1995; FTC, 1998)

IMPLICATIONS FOR FURTHER RESEARCH

The model in Figure 2 suggests a new perspective on conducting research on knowledge management. First, the model suggests that the elements of a company's KMS can be examined not simply as a set of tools, but as fulfilling five design parameters: 1) degree to which the KMS tools are integrated into the work processes and prioritize work over reuse, 2) whether contextual information-gathering is performed in a way that attends to trust and user time, 3) whether personalization increases with use, 4) whether search tools look for strategic rather than contextual similarity, and 5) whether there are human facilitators and review panels to ensure that personalization services, rules, and usages are being exercised as expected and desired. Companies that use a variety of knowledge management services to support their EKP workers – both computerized and not (e.g., LotusNotes, discussion forums, intelligent search, seminars) could be examined along these five design parameters. Thus, this model suggests a framework for comparing KMS initiatives across firms using a much richer framework that simply comparing tools. In addition, the model in Figure 2 offers a theory specific enough to be tested. The general proposition from our theory is that KMSs in firms demonstrating each of these design parameters will have employees innovating faster and with higher quality solutions than firms with KMSs lacking in these design parameters.

In addition to the general research implication of evaluating KMSs according to the five design parameters, each hypothesis suggests further research questions to explore. For example, the 2nd hypothesis suggests improving personalization over time with increased use. However, with EKPs, there is such a variety of contexts, problems, and individuals involved in the process that repeated use may not lead to the gains in personalization that derive from experience. Therefore, it may be that certain types of experiences are more important than others to demonstrably show personalization over time. For example, it may be that knowledge injection opportunities are particularly helpful for problems that are better defined than for less well-defined problems. In this case, users might be encouraged to initially use the services particularly with well-defined problems before trying it out on ill-defined problems. Or, the services can provide increasing personalization more quickly for novices in a discipline than experts (because the uniqueness of the problems expert solve are difficult to generalize), and thus targeting novices initially may lead to greater use and problem-solving performance.

There are many limitations of this framework (which will be worked through as we proceed with our action research at our partner organization and which we can report on in the March workshop). First, our definition of a knowledge delivery service needs to be strengthened. At one extreme, everything is a knowledge delivery service in an EKP since EKPs are, by definition, focused on knowledge delivery and development. At the other extreme, few organizations today have knowledge delivery services as we use the term here since existing services in organizations today are less focused on injecting new knowledge and more focused on simply passively providing knowledge when requested. Thus, clarifying the boundaries around our main construct is necessary. A second limitation is that our definition of personalization is not yet concise enough. At the one extreme, any effort to customize a response to a query (so any search of google) is a personalized service; yet few companies today offer the type of personalized services we are interested in: where the personalization is based not simply on what the user typed in, but based on the user's context (past and present) in which the typing was done. Clarifying our boundaries around the term 'personalization' is needed. Finally, we need yet to operationalize these concepts into measurable yet generalizable characteristics of knowledge services. We have to date operationalized them in very specific terms, such

as specific services (e.g., Legacy reviews, Technical Questions, Discovery Support Service). Creating more generalizable operationalizations for measurement is our next step.

IMPLICATIONS FOR PRACTICE

It is perhaps too early in our conceptualization to discuss implications for practice. However, in working with our partner firm, we have discovered that our framework has helped to identify thus far a few practical implications. These include:

- Any KMS should use the opportunity when an individual interacts with it to collect information about the user's work and problem context. Thus, when an individual types in a search query, the user can be asked to provide brief descriptions of the problem and work context, as well as the degree of satisfaction with the query results. As this information is collected over time, correlations can eventually be computed to determine what combinations of queries, problem & work contexts lead to the most satisfying results.

- Although EKP workers place high value on individualism, there are likely to be lots of similarities based on roles (e.g., project managers generally need to end meetings with action items, mechanical specialists generally worry about a product's hardware). Companies can begin to integrate knowledge delivery into how work is performed by redefining job descriptions not by tasks people perform but by the knowledge they need. Then, baselines for user profiles can be developed to correspond to these roles, jump-starting the personalization of the knowledge delivery services.

- An idealized view of a set of completely computer-based personalized knowledge delivery services cannot yet be supported either by the social system nor technology. Assuming an absolutely trustworthy environment, voice recognition is not sufficiently accurate, people are unlikely to feel comfortable wearing a microphone, and the amount of storage space required to gather and store information generated by each employee each minute of the day is untenable. Thus, knowledge delivery services need to be targeted. Our framework suggests targeting them to the highly innovative work of users since, we argue, the payoff of unplanned innovation can result. Moreover, we suggest that, by conceptualizing KMS as services rather than a system, any organization can start small by piloting individual services and observing the outcomes. Eventually, with continued commitment to the piloting process, a portfolio of services will have been created.

- By considering knowledge delivery services as a combination of human and computer-based knowledge service providers, organizations can begin to develop

new role descriptions. For example, one implication of our framework is to train people in a variety of different knowledge delivery services, just as firms train individuals in problem-solving techniques. For example, individuals need to be trained in opportunity-recognition (i.e., the ability to identify opportunities for knowledge injection), a variety of different knowledge sources, use of low-probability search engines, knowledge capture, and the need to maintain an absolutely trustworthy environment. Thus, if a design team is about to embark on a design meeting, there ought to be a service available where they ask a human facilitator to be present to capture the design rationale at the meeting. In addition, we have suggested that human facilitators need to be supported by computer-based tools. This suggested that services delivered to facilitators might be different than services delivered to the users themselves.

In sum, we argue that, just as our conceptualization introduces new research questions into the research domain of knowledge management, our conceptualization encourages organizations to venture into new areas of knowledge management practice.

REFERENCES

Alba, J., Lynch, J., Weitz, B., Janiszewski, C., Lutz, R., Sawyer, A. and Wood, S. 1997. "Interactive Home Shopping: Consumer, Retailer, and Manufacturer Incentives to Participate in Electronic Marketplaces," Journal of Marketing, 61, pp. 38-53.

Alavi, M. & Leidner, D.E. 2001 Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. MISQ, 25(1), 105-136.

Allen, T.J. 1977. Managing the Flow of Technology: Technology Transfer and the Dissemination of Technological Information Within the R&D Organization. Cambridge, MA: MIT Press.

Ancona, D.G. and Caldwell, D.F. 1990. Information technology and work groups: The case of new product teams. In J. Galegher, R.E. Kraut, & C. Egido (Eds) Intellectual Teamwork: Social and Technological Foundations of Cooperative Work. Hillsdale, NJ: Lawrence Erlbaum.

Argote, L. 1999 Organizational learning: Creating, retaining and transferring knowledge. Norwell, MA: Kluwer

Balachandran, K., Buzydlowski, J., Dworman, G, Kimbrough, S.O., Shafer, T., and Vachula, W. (1999), MOTC: An Interactive Aid for Multidimensional Hypothesis Generation, Journal of Management Information Systems, 16 (1), pp. 17-36.

Bitner, M.J., Booms, B.H. and Mohr, L.A. 1994. Critical service encounters: the employee's viewpoint," Journal of Marketing, 58, pp. 95-106.

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

Blair, D. and Kimbrough, S. O. 2002. Exemplary Documents: A Foundation for Information Retrieval Design, Information Processing and Management, 38, 363-379.

Brown, J.S. & Duguid, P. 1991 Knowledge and Organization: A Social-Practice Perspective. Organization Science, March-April, 198-213

Brown, J.S. & Duguid, P. Organizational learning and communities of practice: Toward a unified view of working, learning, and innovation. Organization Science, 2001, 40-57.

Chellappa, R.K. & Sin, R. 2002 Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma, Proceedings of the 7th Annual Conference on Information Systems and Technology, INFORMS, San Jose, CA, 2002, pp. -.

Cingil, I., Dogac, A. & Azgin, A. 2000. A Broader Approach to Personalization, Communications of the ACM, 43 (8), 136-141.

Clay, K., Krishnan, R. and Wolff, E. 2001. Prices and Price Dispersion on the Web: Evidence from the Online Book Industry, Journal of Industrial Economics 49 (4) 521-539.

Conklin, E.J. & Yakemovich, K.C., B. 1991 A process-oriented approach to design rationale. Human-Computer Interaction, 6(3&4), 357-391.

Cooper, L.P. in press. A research agenda to reduce risk in new product development through knowledge management: A practitioner perspective. Journal of Engineering and Technology Management.

Couger, J.D. 1996. Creativity and innovation in Information Systems Organizations. Thomson Publ.

Culnan, M.J.1995. Consumer Awareness of Name Removal Procedures: Implications for Direct Marketing, Journal of Direct Marketing 9(2), 1995, 10-19.

Culnan, M.J. 2000. Protecting privacy online: Is self-regulation working? Journal of Public Policy & Marketing 19 (1), 20-26.

Davenport, T.H., De Long, D.W., & Beers, M.C. 1998. Successful knowledge management projects. Sloan Management Review, Winter, 43-56

DeSanctis, G. & Poole, M. 1994 Capturing the complexity in advanced technology use: Adaptive structuration theory, Organization Science, 5(2), 121-147.

Derlega, V., Metts, S., Petronio, S. and Margulis, S. Self-Disclosure, 1993.

Doney, P.M. and Cannon, J.P. 1997. An Examination of the Nature of Trust in Buyer-Seller Relationships, Journal of Marketing, 61, 35-51.

Dworman, G., Kimbrough, S.O., and Patch, C. 2000, On Pattern-Direct Search of Archives and Collections, Journal of the American Society for Information Science 51, no. 1.

Faniel, I. and Majchrzak, A. 2002 An exploratory study of the factors associated with successful reuse of other's knowledge. Association for Information Systems Conference, Dallas, Tx. August.

FTC, F.T.C. 1998 Fraud Could Slow Growth of Electronic Commerce. FTC Press Release FTC File No. P97-4406, Federal Trade Commission, June 25 1998.

Fung, R. & Lee, M. 1999. EC-Trust "(Trust in E-commerce): Exploring the Antecedent Factors, Proceedings of the 5th Americas Conference on Information Systems, 517-519.

Ganesan, S. "Determinants of Long-Term Orientation in Buyer-Seller Relationships," Journal of Marketing (58), 1994, pp. 1-19.

Gordon, M.D. & Lindsay, R. K. 1996., Toward Discovery Support Systems: A Replication, Re-examination, and Extension of Swanson's Work on Literature Based Discovery of a Connection Between Raynaud's and Fish Oil, Journal of the American Society for Information Science, 47(2):116-128.

Grant, R.M. 1996. Prospering in dynamically competitive environments: Organizational capabilities as knowledge integration. Organization Science, 7(4), 375-387.

Hansen, M.T., Nohria, N. & Tierney, T. 1999. What's your strategy for managing knowledge? Harvard Business Review, March-April, 106-116.

Hargadon, A. & Sutton, R.I. 1997. Technology brokering and innovation in a product development firm. Administrative Science Quarterly, 42, 716-749.

Hollingshead, A, Fulk, J, & Monge, P. 2002. Fostering internet knowledge-sharing. In P. Hinds & S. Kiesler (Ed) Distributed Work. Cambridge, MA: MIT Press.

Jarvenpaa, S.L. & Staples, D.S. 2000. The use of collaborative electronic media for information sharing: an exploratory study of determinants Journal of Strategic Information Systems, 9, 129-154.

Koller, M. 1988. Risk as a determinant of trust, Basic-and-Applied-Social-Psychology (9:4), 265-276.

Kramer, J., Noronha, S. and Vergo, J. 2000. A User-Centered Design Approach to Personalization, Communications of the ACM (43:8), 45-48.

Leifer, R., McDermott, C., O'Connor, G., Peters, L., Rice, M., and Veryzer, R. (2000) Radical Innovation. Boston, MA: Harvard Business School Press.
Lind, A. and Tyler, T. 1988. The Social Psychology of Procedural Justice, Plenum Press, New York, NY,.

Majchrzak, A., Cooper, L.P. & Neece, O.E. (in press) Knowledge reuse in the radical innovation process at the Jet Propulsion Laboratory. In J. Holm & J. Liebowitz (Ed) Knowledge Management Technologies and Applications at NASA.

Majchrzak, A., Rice, R.E., Malhotra, A., King, N. & Ba, S. 2000. Technology Adaptation: The case of a computer-supported inter-organizational virtual team. MISQ, 24(4), 569-600.

Malhotra, A., Majchrzak, A., Carman, R., and Lott, V. 2001. Radical innovation without collocation: A case study at Boeing-Rocketdyne, MISQ, 25(2), 229-249

Markus, M.L. 2001. Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success. Journal of Management Information Systems, 18(1), 57-93.

Markus, M.L., Majchrzak, A. and Gasser, L. 2002 A design theory for systems that support emergent knowledge processes" MISQ, 26(3), 179-212

Mayer, R.C., Davis, J.H. and Schoorman, F.D. 1995. An integration model of organizational trust, The Academy of Management Review (20:3), 709.

Mittal, B. & Lassar, W.M.1996. The Role of Personalization in Service Encounters, Journal of Retailing (72:1), 95-109.

Mobasher, B., Cooley, R. & Srivastava, J.2000. Automatic Personalization Based on Web Usage Mining, Communications of the ACM (43:8), 142-151.

Mulvenna, M.D., Anand, A.S. & Buchner, A.G. 2000. Personalization on the Net using Web Mining, Communications of the ACM 43 (8), 123-125.

Nidumolu, S.R., Subramani, M., & Aldrich, A. 2001. Situated learning and the situated knowledge web: Exploring the ground beneath knowledge management. Journal of Management Information Systems, 18 (1), 115-150.

Orr, J. 1996. Talking about machines: An Ethnography of a Modern Job. IRL Press, Ithaca, NY

Quilter, L. 2002. The continuing expansion of cyberspace trespass to chattels, Berkeley Technology Law Journal 17 (1),. 421-443.

Schultze, U. & Boland, R.J. 2000 Knowledge management technology and the reproduction of knowledge work practices. Journal of Strategic Information Systems, 9, 193-212.

Schurr, P.H. and Ozanne, J.L. 1985. Influences on Exchange Processes: Buyers' Preconceptions of a Seller's Trustworthiness and Bargaining Toughness, Journal of Consumer Research 11 (4), 939-953.

Selvin, A.M. (in press) Fostering collective intelligence: Helping groups use visualized argumentation. In P.A. Kirschner, S. Buckingham Shum, & C. S. Carr (Ed) Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making. Springer-Verlag: London

Sheehan, K.B. and Hoy, M.G. 2000. Dimensions of privacy concern among online consumers, Journal of Public Policy & Marketing 19 (1) 62-72.

Singh, J. & Sirdeshmukh, D. 2000. Agency and trust mechanisms in consumer satisfaction and loyalty judgments, Academy of Marketing Science Journal 28 (1), 150-167.

Smith, J.H., Milberg, S.J. & Burke, S.J. 1996. Information Privacy: Measuring Individuals' Concerns About Corporate Practices MIS Quarterly 20 (2), 167-196.

Solomon, M.R., Surprenant, C.F., Czepiel, J.A. & Gutman, E.G. 1985. A Role Theory Perspective on Dyadic Interactions: The Service Encounter, Journal of Marketing 49(4), 99-111.

Stone, E.F., Gardner, D.G., Gueutal, H.G. & McCLure, 1983. A Field Experiment Comparing Information-Privacy Values, Beliefs, and Attitudes Across Several Types of Organizations, Journal of Applied Psychology, 68 (3), 459-468.

Stone, E.F. & Stone, D.L. 1990. Privacy in Organizations: Theoretical Issues, Research Findings, and Protection Mechanisms, In Research in Personnel and Human Resources Management, K. M. Rowland and G. R. Ferris (Ed.), 8, JAI Press, Greenwich, CT, 349-411.

Swanson, D.R. 1988, Migraine and Magnesium: Eleven Neglected Connections. Perspectives in Biology and Medicine 31(4), 526-557.

Szulanski, G. 2000. The process of knowledge transfer: A diachronic analysis of stickiness. Organizational Behavior and Human Decision Processes, 82(1), 9-27.

Volokh, E.2000 Personalization and Privacy, Communications of the ACM , 43(8), 84-88

Zucker, L.G. 1986 Production of Trust: Institutional Sources of Economic Structure, 1840-1920, Research in Organizational Behavior (8), 53-111.