Journal of Strategic Information Systems 9 (2000) 175–191

Strategic Information Systems

www.elsevier.com/locate/jsis

The effects of knowledge management systems on emergent teams: towards a research model

Peter H. Gray

School of Business, Queen's University, Kingston, Ontario, Canada K7L 4S8

Abstract

This article describes how knowledge management systems can enhance the effectiveness of teams that analyze complex, non-recurring problems by improving the way that team composition evolves. Knowledge management systems reduce the costs of searching for specialized knowledge resources, making it more likely that teams will incorporate a diversity of knowledge. Drawing on the concept of requisite variety, this article argues that increases in team knowledge variety lead to improvements in the effectiveness of the solutions generated by a team, which in turn enhance their organization's adaptive ability. This process also reinforces the existing distribution of knowledge within the organization, increasing employee specialization. The author develops a series of propositions and combines them into a research model from which he draws implications for researchers and managers.

Keywords: Knowledge management systems; Requisite variety; Knowledge maps

> *People who ran those factories in the brute-force economy of the past liked large numbers of predictable, interchangeable, don't-ask-why workers for their assembly lines…. Third Wave economies, by contrast, will require (and tend to reward) a radically different kind of worker-one who thinks, questions, innovates, and takes entrepreneurial risk. Workers who are not easily interchangeable. Put differently, it will favor individuality.*
>
> (Alvin and Heidi Toffler, 1995)

1. Introduction

Knowledge management systems (KMS) have been the subject of considerable interest

E-mail address: pgray@business.queensu.ca (Peter H. Gray).

0963-8687/00/$ - see front matter
PII: S0963-8687(00)00040-8

by academics and practitioners over the past decade, yet little cumulative empirical research has been conducted to establish the causal mechanisms by which KMS can influence organizational performance. This article develops a series of testable propositions to describe the connection between KMS use and the quality of solutions produced by "emergent" teams (characterized by an ongoing evolution in team composition that does not follow a predefined pattern). In addition, it proposes that ongoing KMS use increases employee specialization, which in turn reinforces KMS use. By putting forward a researchable model of the effects of KMS use on organizational adaptation and employee specialization, this article calls attention to a set of testable intermediate effects that may help advance research in this emerging field.

KMS are a class of information systems that focus on creating, gathering, organizing and disseminating an organization's knowledge (Alavi and Leidner, 1999). Two common types of KMS are *knowledge repositories* and *knowledge maps*. Broadly speaking, the former are databases of documents written by knowledgeable individuals, while the latter are searchable indexes or catalogues of expertise held by individual employees (Davenport and Prusak, 1998; Crowe, 1997; Ruggles, 1998).

The proposed effects of KMS use can be summarized as follows. Both knowledge repositories and knowledge maps improve employees' ability to search for and find knowledgeable individuals within an organization. By making knowledge searches more effective, KMS improve the variety of knowledge present on emergent problem-solving teams. Improved team knowledge diversity leads to more accurate and complete analysis of complex problems, improving the effectiveness of the solutions teams generate and, ultimately, enhancing an organization's ability to adapt to its environment. Successful KMS use improves the perceived value of KMS, which leads to higher levels of use. Finally, individuals who are regularly asked to contribute specific knowledge will become increasingly specialized, which will increase the need for integrating mechanisms such as KMS. Establishing a causal model of the effects of KMS remains a major research issue that needs to be addressed; these propositions are a first step towards such a model.

The remainder of this article is organized in five sections. The first discusses the use of emergent teams to improve organizational adaptation. The second introduces the principle of requisite variety and extends it to apply to knowledge diversity among team members. The third explores the process of assembling teams with a sufficient variety of knowledge to solve complex, non-recurring problems. This is followed by a discussion of the manner in which KMS can improve the variety of knowledge represented on a team, and the likely results of improved variety on solution effectiveness, organization adaptability and employee specialization. The final section presents research propositions that flow from this analysis as well as implications for managers.

2. Using teams to enhance organizational adaptation

Organizations increasingly operate in complex and uncertain environments. To maintain their fitness for survival, they must continuously adapt to their evolving circumstances (Van Valen, 1973). Purposeful adaptation begins with problem-sensing behavior, described by Kiesler and Sproull (1982) as noticing, interpreting and incorporating stimuli. Effective adaptation in complex and rapidly evolving environments necessitates additional effort devoted towards analysis and interpretation.

Problems drawn from rapidly changing environments require non-programmed decisions (Simon, 1977) that are novel, unstructured and lacking in clear precedent. Such decisions involve non-routine processes with few set procedures and a high degree of uncertainty (Van de Ven et al., 1976), and require the application of knowledge that is not (and cannot be) known by a single individual (Tsoukas, 1996). Ad hoc groups often address complex, non-routine issues emerging from rapidly changing environments (Finholt et al., 1990). Because it is a communication-intensive form of knowledge integration, problem solving by such teams is typically reserved for unusual, complex and important tasks (Grant, 1996).

Assembling an appropriate problem-solving team presents a significant managerial challenge when the full scope of a complex problem is not evident at the outset. With only a partial sense of the complexity of a problem, a manager is unlikely to identify all of the specific knowledge a team will require to produce a solution. Teams that lack sufficient diversity in knowledge to comprehend all aspects of a complex problem will be hampered in their ability to analyze the problem and solve it. As the multifaceted nature of a complex problem emerges, teams must therefore evolve to incorporate new members with complementary knowledge.

KMS can help team members find individuals with particular knowledge to help analyze complex problems. KMS are clearly not the only way for a team to search for new members; however, when compared to manual searches, searches using KMS require less effort. This improves the likelihood that a team will increase its diversity of knowledge to analyze a problem. The following section addresses issues surrounding the assembly of sufficient knowledge variety on problem solving teams, beginning with an introduction of the principle of requisite variety.

3. The principle of requisite variety

Drawing from cybernetics theory Ashby (1968), developed the “law” (or principle) of requisite variety. Ashby demonstrated that a control system cannot always respond optimally to environmental stimuli when the variety in the stimuli is greater than the variety in the system’s internal set of responses. The following example helps draw out important implications of this principle.

A simple (and oft-used) illustration of a control system is a thermostat connected to a furnace. When the house temperature drops below a certain level, the thermostat activates the furnace until the temperature rises to the desired level. Consider for a moment a simple programmable thermostat that features a 24-hour clock; the homeowner can set a combination of times and desired temperatures, and the thermostat controls the furnace to maintain those temperatures. The thermostat seems at first to be well suited to its environment, as it can, for instance, turn the temperature up at 6 a.m. and down again at 9:30 p.m.

Now, consider an increased level of variety in the environment: weekends, when the homeowner may sleep in or stay up late. The simple thermostat cannot react optimally to the variety in the environment because it lacks an internal representation of a seven-day cycle. To properly match this variety, the thermostat requires an internal calendar to track the current day of the week and apply a different set of rules for each day.

Ashby's principle does not imply that a control system will always respond sub-optimally when variety increases, merely that it cannot always respond optimally. The thermostat described above will respond optimally for much of the week, but it cannot always do so when there are cycles to the homeowner's heating preferences. Thus, a prerequisite for any effective control system is that its set of internal states must match or exceed the number of possible states present in its external environment.

Organizations can be seen as control systems: they undertake actions that influence their environment, which produces feedback that influences future actions. Organizations are obviously much more complex than thermostats — they face far fewer clearly delineated problems featuring easily circumscribed choices. Regardless, the principle of requisite variety has clear implications for organizations: to respond optimally, the organization must be capable of forming representations that match the complexity present in the external environment. Galunic and Eisenhardt (1994) use contingency theory to come to the same conclusion, arguing that to ensure survival, organizations must respond to external complexity with greater internal complexity.

In addition to being proven mathematically by Ashby, the principle of requisite variety has been widely used in systems theory and organization theory. For example, Aulin (1982) extended it to develop the *law of requisite hierarchy*: the larger the variety in the environmental perturbations that need to be compensated for, the larger the functional hierarchy of control levels required in living systems. Beer (1974) extended Ashby's work, contending that a system that insulates itself from environmental variety will become highly unstable. Beer argued that organizations must therefore absorb environmental variety by either amplifying internal variety or attenuating environmental variety.

Ashby's principle is based on two assumptions that do not necessarily apply to individuals dealing with the complex nature of their organization's external environment. The first assumption is that environmental states manifest themselves in such a way that they are clearly visible to individual decision makers. The second is that, after observing the environmental state, decision makers will know how to respond optimally. Neither assumption is always valid in an organizational context; in fact, both are often invalid. Relaxing these assumptions yields new corollaries that more accurately reflect the human element in organizations, as explored below.

3.1. Information and requisite variety

The principle of requisite variety can be extended to encompass the role that information plays as organizations respond to changes in the external environment (Choo, 1995). By relaxing the assumption that changes in environmental states are always clearly visible to the organization, Choo asserts that when organizations fail to detect certain environmental stimuli, they may not form accurate representations of the environment. Inaccurate representations (based on faulty or missing information) can therefore lead to suboptimal responses. This information-processing perspective maintains that a necessary (but not sufficient) prerequisite for an effective organization is that its variety of *information* *sources* must match the variety present in its external environment (e.g. Choo, 1995; Galbraith, 1973; Simon, 1945; Tushman and Nadler, 1978).

This assertion rests on the premise that individuals' ability to make effective decisions for the organization is limited primarily by the quality and variety of information available to them. Indeed, organizations may react ineffectively to changes in their environment if they do not possess sufficient information about those changes. Yet, as Aadne et al. (1996) argue, this is a limited perspective on decision making and problem solving; because information requires interpretation, the critical failure of organizations to react effectively to changes in the environment is more likely to stem from their inability to properly interpret environmental change. This suggests that there are limits to Ashby's implicit assumption that individuals will know how to respond optimally to changes in environmental states as long as they can form sufficiently detailed representations of their environment.

3.2. Knowledge and requisite variety

The major difficulty with focusing on the role of information when applying the principle of requisite variety lies in the definition of the word *information*. A key assumption tied to the information-processing view of the firm is that information contains meaning that is independent of the person who receives it. However, the same information does not necessarily have the same meaning for different recipients (Argyris, 1980). What is information to some is not to others; converting data into information requires knowledge (Drucker, 1988). As Tricker (1992) argues, the word *information* implies the process of becoming informed. Such a process is inherently subjective: if an individual gains new understanding after receiving data, then he or she is informed by it. Another person may receive the same data and not understand its meaning; that person would not be informed by that data.

The ability to grasp the meaning inherent in data depends on an individual's knowledge of the subject domain. The greater the understanding an individual has of a subject domain, the better he or she will be able to be informed by data drawn from that domain (Cohen and Levinthal, 1990). Thus, additional problems arise when an organization lacks the knowledge to interpret the data, make sense of it and draw conclusions from it. It follows that a necessary (but not sufficient) prerequisite for an effective organization is that it must be capable of *understanding* the complex phenomena in its environment. That is, organizations that lack a sufficient diversity of knowledge to understand the stimuli present in the external environment cannot always react optimally to changes in the environment.

A simple example of a complex environmental stimulus is a change in market conditions that suggests the need to develop a new product or service; the firm that responds to this challenge by forming a team staffed solely by individuals drawn from one functional area is unlikely to discover all the key requirements that a successful product would have. A multifunctional team that brings a wide variety of knowledge to bear on the problem is more likely to discover and therefore meet the complex requirements of a successful new product/service. The integration of diverse knowledge held by a variety of people is discussed in detail in the following section.

4. Team formation and knowledge variety

Organizations increasingly use teams to solve complex problems (Lawler et al., 1995); teams are a mechanism for pooling and using the diverse knowledge and skills of employees to accomplish mutual goals (Drucker, 1994). To be effective, teams require sufficiently diverse knowledge to properly assess and understand the problems they face. Inadequate problem assessment can lead to poor decisions, inferior products and unreasonably high costs associated with searching for information and evaluating solutions (Tushman and Nadler, 1978).

The full scope of a complex problem is rarely evident when the problem is first encountered. A manager facing such a problem must therefore assemble a team without fully understanding the problem domain. Operating under bounded rationality (Simon, 1991) and with incomplete information, it follows then that a manager may not properly identify all the required knowledge bases when putting together the team. This may result in an initial mismatch of variety: the system (the team) does not have the variety required to understand the phenomena present in the external environment (the problem).

Inadequate knowledge variety at inception does not necessarily doom a team to failure. When a team is aware of its shortcomings, it can search for and incorporate the missing knowledge into its structure by adding new members. Indeed, with a sufficiently complex problem, there may be no final structure for such emergent teams; individuals who are required only for a particular phase of analysis or design contribute their specialized expertise and then move on.

This line of reasoning substantiates Drucker's (1988) predictions of a new organizational form characterized by work performed by specialists brought together in task forces that cut across traditional organizational lines. Drucker sees managers as deciding team composition, and others have also focused on top-down (albeit iterative) team design (for instance Mankin et al., 1998). Yet, when managers know only a fraction of what their subordinates know, managers are inefficient co-ordinators (Grant, 1996). Thus, the particular configuration of knowledge represented on a team is likely to emerge over time, and team members may be the best people to select new members as they have the best understanding of the problem and the specialized knowledge required to solve it.

4.1. Perceived costs and benefits of knowledge searches

The term *knowledge search* in this article refers to the process of trying to find an individual who has specialized knowledge relevant to a particular problem.[1] Knowledge searches require time and effort to identify potential knowledge holders and contact them to assess the relevance of their knowledge. Davenport and Prusak (1998, pp. 88–89) describe the suboptimal effect of human nature on this process, arguing that satisficing behavior on the part of individuals limits their willingness to conduct a thorough search for

[1] This term has been used elsewhere to refer to the process of searching for documents containing codified knowledge. While a considerably body of research deals with searching the contents of electronic databases, relatively little research addresses the use of such systems to find individual experts. The effects proposed in this article are not expected to apply to searches for codified knowledge stored in documents.

the right expertise. This increases the probability that decisions will be made using incomplete or deficient knowledge.

According to the cost-benefit principle of cognitive effort (Payne, 1982), decision makers select a strategy by considering the amount of cognitive effort it demands and the benefits it will yield (Benbasat, 1989). In the context of knowledge searches, the costs are the perceived time and effort required to locate and contact an individual and assess whether he or she possesses the required knowledge, and the benefits are the perceived likelihood that a contacted individual will possess the required specialized knowledge. Thus, someone who is searching for specialized knowledge is likely to contact the person with whom they associate the largest ratio of benefit to cost. Should that person not have the required knowledge, then each subsequent person contacted would represent a poorer ratio of benefits to costs. This process presumably continues until the cost of contacting one additional person outweighs the benefits. The searcher gives up the search, preferring perhaps to reframe the problem in a simplified manner or to rely on less appropriate (but more easily accessible) knowledge bases. In either case, the specialized knowledge required to match the complexity of the task is not applied. The searcher has acted rationally, yet from the organization's perspective a less-than-ideal solution is the likely result.

5. Use of KMS in knowledge searches

Traditional ways of searching for knowledge become more difficult as an organization becomes larger and more distributed and as the knowledge of its members becomes more specialized. This in turn increases the likelihood of premature knowledge search termination. KMS can reduce the likelihood of premature search termination by decreasing the amount of effort required to locate likely sources of expertise, thereby improving the cost-benefit trade-off associated with surveying the knowledge of a wide range of people (see Fig. 1). KMS offer two new ways to locate individual experts. First, teams can find individuals with the sought-after knowledge by consulting a knowledge map that lists employees and their areas of expertise. Second, employees can be identified by tracing the authors of key documents on the subject in question in a knowledge repository.

In Fig. 1, the term *manual search* refers to traditional ways of finding knowledge (finding the right people and then talking to them), whereas *technology-assisted search* indicates the use of KMS to locate knowledge. Note that Fig. 1 presumes that all other things are held constant: organizational culture, individual roles, reward systems, political issues and a wide variety of other factors will greatly influence knowledge sharing, but are beyond the scope of this model. It is also assumed that the searcher has some experience using the technology; for those who have none, the technology-assisted search curve may initially be significantly higher.

Fig. 1 depicts the relationship between the costs and benefits of knowledge searches. For manual searches, the marginal cost of talking to one extra person is small at first (ask the person in the office next to you) and grows rapidly (try to find someone in a different department who might know something you need). The costs associated with a technology-assisted search may be initially higher than for a manual search (it is easier to

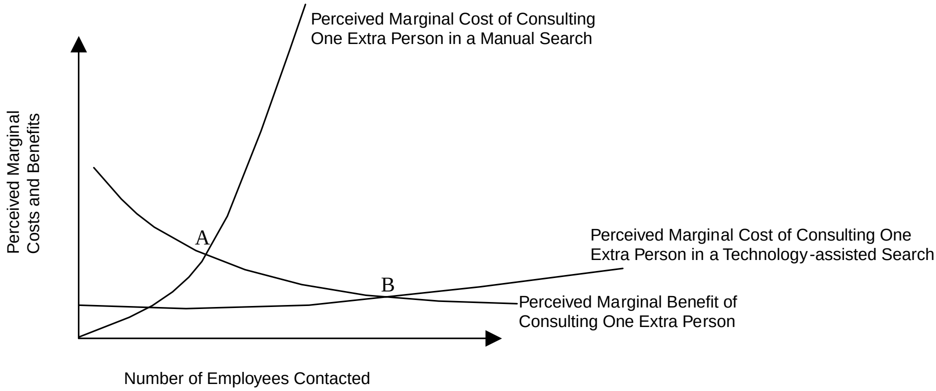

Fig. 1. Proposed relationship between manual and technology-assisted searches.

ask the person next to you than to search through KMS). However, such costs grow at a much slower rate, as the effort expended in searching electronically to find someone with relevant knowledge increases almost imperceptibly as the number of people referenced in the system grows. In this way, KMS improve an individual's ability to perform ad hoc searches for individuals with specialized knowledge.

When an organization is so large and distributed that people with relevant knowledge may not be easy to reach, some individuals searching manually will give up before finding the knowledge they need. This corresponds to point A in Fig. 1, beyond which the perceived search cost exceeds the benefit of contacting one more person. However, those who use KMS in their search are able to perform much broader searches for individuals with suitable knowledge before concluding that it is not worth continuing (point B in Fig. 1). This suggests the following proposition (see Fig. 3 for a diagram of all propositions).

P_1: KMS use increases the efficiency of knowledge searches in large, distributed organizations

The question of how large and distributed an organization must be to benefit from KMS in knowledge searches remains open to speculation, as it has not yet been demonstrated empirically. However, an increase in efficiency is not expected when organizations are small enough that most employees' contact networks overlap significantly and when those employees work in the same location. Likewise, experienced employees may be able to effectively search much larger organizations than can new employees.

Improvements in the efficiency of knowledge searches are likely to have a beneficial effect on team knowledge variety; if it is easier to search, people will be more likely to do so and thus more likely to locate the desired sources of expertise. With all else held constant, the improved ability to locate potential knowledge contributors improves the variety of knowledge accessible to a team, as summarized in Proposition 2:

P_2: more efficient knowledge searches increase the variety of knowledge represented on a team

The ability of a team to enlist new members is, of course, limited by their availability, their interest in participating and the incentives for doing so. The ability to change team composition is also constrained by organizational constructs such as culture, flexibility, hierarchy and slack, as well as individual factors such as motivation, disposition, attitudes towards the team, and pre-existing commitments.

Based on the previous discussion of requisite variety, increases in team knowledge variety should have a positive effect on problem analysis (subject to some limitations, discussed below). This follows Blackburn's (1971) principle of complementarity: full understanding of a phenomenon requires the application of a number of complementary theories or descriptions. Though such theories or descriptions are inherently incomplete and possibly overlapping, they cannot be subsumed into each other. A full description of phenomena is therefore only possible through the application of many diverse knowledge bases; the more complex the phenomena, the more variety in knowledge required. Increased knowledge variety is therefore expected to lead to more complete problem analysis, as has been demonstrated extensively in product development literature (e.g. Cooper, 1994). Proposition 3 is, therefore, as follows:

P_3: increased team knowledge variety leads to more effective problem analysis

Naturally, there is a limit beyond which further knowledge variety is pointless. However, emergent teams that seek out new knowledge as they uncover new facets of a complex problem are self-governing; they increase variety in response to characteristics of the problem, and are unlikely to expand variety beyond what is helpful to the task.

There is also a point at which increased variety may needlessly complicate group processes. Groups become less effective as the amount of effort required to coordinate members and activities grows (e.g. Brooks, 1982). Increases in knowledge diversity may also reduce members' ability to communicate; without common points of reference, individuals with diverse knowledge may be unable to collaborate (Cohen and Levinthal, 1990).

Yet there are clearly situations that feature a favorable balance between the benefits of knowledge diversity and process losses. When dealing with non-routine problems arising from complex and rapidly changing environments, teams are more likely have too little knowledge variety than too much.

If overall effectiveness is a combination of group process effectiveness and analysis effectiveness, then improvements in knowledge diversity will be beneficial only to a certain level (shown as a movement from D to D' in Fig. 2). At this point, further increases in knowledge diversity will generate more process losses than improvements in analysis effectiveness. Though not all problem-solving teams would benefit by improved knowledge variety, this article proposes that teams analyzing complex, novel problems that are early in their formation are likely to be to the left of D' than the right.

Two assumptions are necessary to bridge the connection between completeness of analysis, solution effectiveness and organizational adaptation. The first is that teams are

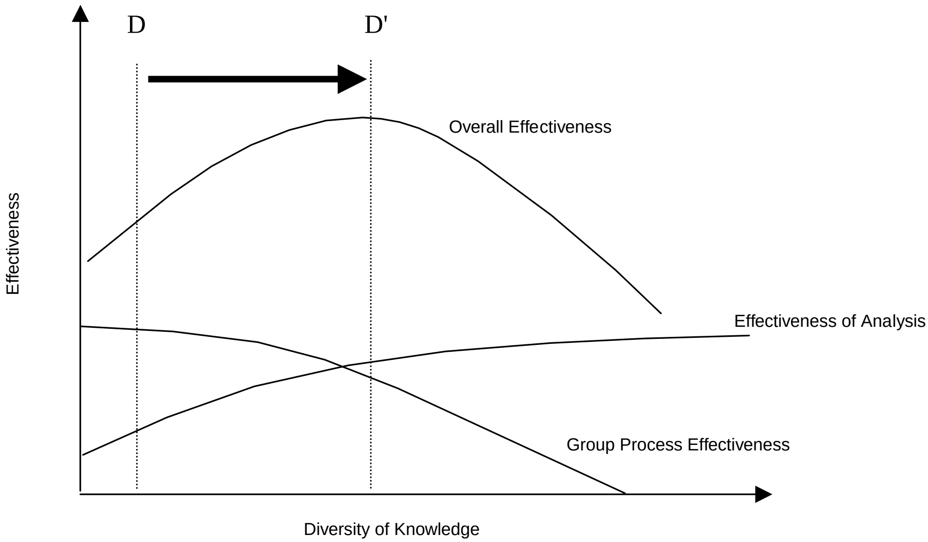

Fig. 2. Process losses and analysis effectiveness.

likely to come up with superior solutions when they fully understand a problem. A number of theorists have made this link (e.g. Simon, 1945; Mintzberg et al., 1976). The weakness of this assumption is implementation: superior analysis may well be wasted if it does not inform the design of a solution, or if the design is poorly implemented. Any number of conditions can interfere with implementation, yet none are expected to differ systematically with KMS use. More complete analyses are therefore assumed to lead to more effective solutions.

The second assumption is that the implementation of superior solutions based on a more accurate understanding of the environment enables the firm to adapt more effectively. Two principle constraints govern the link between solutions and adaptation. First, an organization can reject a solution for any number of reasons; in such a case, and with no direct impact on organizational adaptation, there may still be a residual effect as the knowledge developed in the analysis indirectly improves adaptation by informing future efforts. Second, the environment may evolve in an unexpected direction, making the solution irrelevant or counterproductive. Once again, though the solution does not improve adaptation, the knowledge generated in producing it may improve future problem-solving efforts (Cohen and Levinthal, 1990).

The remaining propositions are ancillary to the adaptation argument but interesting for their effects on knowledge distribution and integration. They take the form of two feedback loops that reinforce KMS use (see Fig. 3).

The first feedback loop centers on employee specialization. Individuals who spend time analyzing problems will improve their understanding of that problem domain; those who were included in a team because of their specialized knowledge are likely to further improve and refine their knowledge in that area. Spending time working on problems within one's area of expertise will tend to reinforce that expertise. When

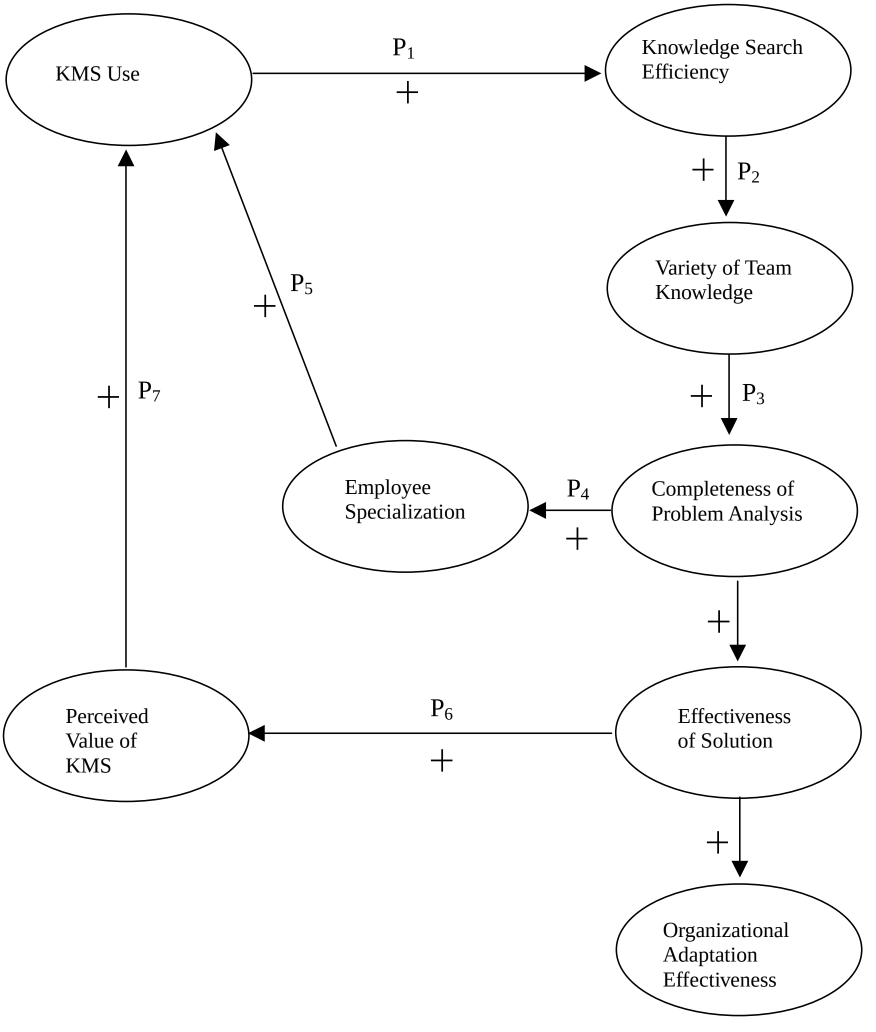

Fig. 3. Proposed effects of KMS use.

compared to an assignment process that does not match individual expertise to problem requirements as effectively, this process tends to make employees increasingly specialized.

P_4: *increased time spent analyzing problems within their area of expertise will increase employees' level of specialization*

This proposition is bounded by an acknowledgement that the process of working on a team has a broadening effect on individuals' perspectives by exposing them to alternative viewpoints. Yet, if the result of KMS use is that individuals are increasingly asked to contribute their specialized knowledge, the result is more likely to be reinforcement and deepening of that knowledge.

This line of reasoning also suggests an increased need for KMS. Increasing employee specialization will heighten the need for some mechanism to locate specialists. This reinforces the conditions under which KMS serve a valuable function to emergent teams. Blau (1970) observed the same effect, noting that increased structural differentiation in employees was associated with higher levels of administrative effort required to integrate employees. KMS act as a substitute for administrative integration of employees; increased levels of employee specialization will therefore lead to more use of KMS to locate that knowledge.

P_5: increased employee knowledge specialization will lead to increased KMS use

The second feedback loop is primarily behavioral. Individuals who participate on teams that have used KMS to improve team knowledge variety, and as a result developed more effective solutions, will tend to look favorably on KMS.

P_6: increased solution effectiveness will improve individuals' perception of the value of KMS

This presumes that the individuals in question are aware that KMS have been effective in improving knowledge diversity. In this model, at least one person used KMS to improve team diversity, and thus at least one person stands to improve their perception of KMS value. Increased perception of value will lead to higher levels of future KMS use.

P_7: increased perception of KMS value will lead to higher levels of use

Fig. 3 depicts these propositions in a single causal chain.

6. Research implications

This article argues that by making it easier to bring a wider variety of specialized knowledge to bear on a complex problem, KMS can improve an organization's understanding of its environment. It illustrates the logic that can form the basis for a stream of research to assess the potential impact of KMS on the quality of problem analysis, and thus the effectiveness of organizational adaptation. In addition, it proposes two feedback loops that reinforce KMS use.

Because the proposed consequences of KMS range from individual to organizational levels of analysis, the full model may not be testable in a single study; researchers may prefer to test smaller groupings of propositions in different ways. P_1 for instance is relatively easy to assess with a controlled experiment or survey. In contrast, P_2 can perhaps be best evaluated by investigating actual KMS operating within organizations. Indeed, organizations intending to implement a KMS provide a rich potential site for an interrupted time-series study (Cook and Campbell, 1979; Cook and Shadish, 1994) that could compare pre- and post-implementation levels of knowledge variety on problem-solving

Table 1
Expected impact of KMS on team knowledge variety

	Problem type	
	Simple, bounded, repetitive	Complex, novel, multifaceted
Duration of Team		
Ongoing	Low	Medium
Project-specific	Low	High

teams. Alternatively, researchers could assess changes in the average amount of time spent by employees in the analysis of complex, non-recurring organizational problems before and after KMS implementation.

P_3 can be assessed by comparing third-party opinions of the completeness of problem analysis across a variety of projects and testing for a relationship between it and an index of team knowledge variety. P_4 may be more difficult to test quantitatively, suggesting a qualitative approach. For instance, longitudinal case studies of organizations using KMS may be better suited assessing changes in levels of employee specialization. Yin (1994, p. 113) proposes a case study method using time-series analysis that would be appropriate for researching the overall impact of KMS on specialization and any resulting effect on KMS use (P_5).

P_6 and P_7 are commonly addressed in research on diffusion of innovation, where the rate of adoption is affected by the user's perceived value of the innovation, which in turn is influenced by the usefulness of the innovation. A wide variety of quantitative and qualitative techniques can substantiate these propositions.

The type of teams that are using KMS will have a significant impact on the likelihood of finding results. The benefits of KMS have been argued in the context of emergent teams, whose membership evolves over time in response to complex, non-routine problems. Other types of teams are not likely to derive the same benefits from KMS use. The expected impact of KMS on various forms of teams is detailed in Table 1. To maximize the likelihood of finding effects, researchers are advised to focus their investigation on settings in which teams are assembled for specific projects dealing with complex, novel and multifaceted/multidisciplinary problems, as such teams are more likely to exhibit an emergent structure.

6.1. Management implications

Three important implications for managers follow from the argument put forth in this article.

The first addresses the connection between organizational flexibility and the likelihood of achieving the proposed benefits of KMS use. Emergent teams are less likely to exist in rigid, bureaucratic organizations than in flexible, organic ones. Although the proposed benefits of KMS set forth in this article are by no means the only ones possible, it seems clear that managers who are hoping to secure them must ensure that employees have sufficient flexibility in their roles to respond to requests for assistance when they arise.

KMS stand to have a stronger impact on improving the composition of emergent teams than of static teams.

The second implication is relevant to both organizations that use KMS and those that do not. Regardless of the role of such systems in team formation, the core of the argument presented above underscores the value of flexible team structures that adapt as team members improve their understanding of the problem at hand. Organizations that currently do not support such emergent team structures may be able to improve team performance by ensuring an adaptive, ongoing and member-driven team formation process without implementing a KMS.

Finally, the potential for positive feedback cycles suggests that the KMS use may become a self-reinforcing process. Once the value of such systems is demonstrated, their use is likely to become commonplace to ensure that emergent teams incorporate sufficiently diverse knowledge into problem analysis. Widespread acceptance and use may take time, but is likely under such conditions.

The considerable interest shown by managers in KMS suggests that they see these systems as having the potential to improve organizational effectiveness. Such an interest will, it is hoped, support in-depth field research to substantiate the many ways that organizations can benefit by applying technology to the task of managing knowledge.

6.2. Limitations and conclusions

In addition to the limitations to each propositions discussed above, three general limitations bear noting.

First, the propositions described above address potential improvements to team problem analysis, but do not deal with important issues concerning group processes. For instance, Ancona and Caldwell (1992) have suggested that product development teams composed of individuals from diverse backgrounds may require greater skills in negotiation and conflict resolution than do homogeneous groups. Although Ancona and Caldwell found that increased diversity did have positive effects on some aspects of team performance by bringing more creativity to problem solving, they also found that it reduced team members' ability to work together, which had a negative impact on performance. Other social psychological issues that may affect team effectiveness include pressures for conformity, domination of the group by an individual or sub-group, and suppression of information to maintain group cohesion. A complete analysis of process losses (Steiner, 1972) that could result from increased knowledge variety is beyond the scope of this article.

The issue of the common language required to integrate complex knowledge is equally important (e.g. Nahapiet and Ghoshal, 1998). It has long been recognized that individuals with different specialized knowledge domains can have difficulty establishing a common language when interacting. The argument presented in this article assumes that the time and effort required to establish such a common language and shared points of reference will not significantly hamper team performance. Further research on the process by which individuals establish such a mutual understanding of terms and their meanings is required.

Lastly, this article assumes that KMS content is both accurate and relevant. Without the expenditure of time and effort to keep KMS current, these systems will fail to yield an advantage to organizations (or, indeed, they may be counterproductive). This relates closely to the issue of cost: do the benefits of a KMS outweigh the costs of implementation? By examining costs solely in the context of knowledge searches, this article has ignored a potentially important factor limiting KMS success.

With these limitations in mind, the propositions set forth in this article represent a testable conceptualization of the effects of KMS on team problem solving, organizational adaptive responses and knowledge specialization. Future research that builds on these propositions to assess the impacts of KMS stands to both inform practitioners and advance theory in this newly emerging field.

Acknowledgements

The author would like to thank Yolande Chan, Brent Gallupe, James McKeen, Darren Meister, David Ramsden and Sandy Staples for their insight and advice in preparing this article. Many thanks also to the guest editor and anonymous referees for their helpful comments on earlier versions of this article, and to the Queen's Management Research Centre for Knowledge-Based Enterprises (http://business.queensu.ca/kbe) for their generous support of this project.

References

Aadne, J.H., Von Krogh, G., Roos, J., 1996. Representationism: the traditional approach to cooperative strategies. In: von Krogh, G., Roos, J. (Eds.), Managing Knowledge. Sage, Thousand Oaks, CA, pp. 9–31.

Alavi, M., Leidner, D.E., 1999. Knowledge management systems: issues, challenges, and benefits. Communications of the AIS 1 (February).

Ancona, D.G., Caldwell, D.F., 1992. Demography and design: predictors of new product team performance. Organizational Science 3 (3), 321–341.

Argyris, C., 1980. Some inner contradictions in management information systems. In: Lucas, Land, Lincoln, Supper (Eds.), The Information Systems Environment. North-Holland, Amsterdam, pp. 15–35.

Ashby, W.R., 1968. Variety, constraint, and the law of requisite variety. In: Buckley, W. (Ed.), Modern Systems Research for the Behavioral Scientist. Aldine, Chicago, IL, pp. 129–136.

Aulin, A., 1982. The cybernetic laws of social progress: towards a critical social philosophy and a criticism of Marxism. Pergamon Press, Oxford.

Beer, S., 1974. Designing Freedom. House of Anansi Press, Don Mills, Ontario.

Benbasat, I., 1989. Laboratory experiments in information systems with a focus on individuals: a critical appraisal. In: Benbasat, I. (Ed.), The Information System Research Challenge: Experimental Research Methods. Harvard Business School Press, Boston, MA, pp. 97–111.

Blackburn, T.R., 1971. Sensuous-intellectual complementarity in science. Science 172, 1003–1007.

Blau, P.M., 1970. A Formal theory of differentiation in organizations. American Sociological Review (35), 201–218.

Brooks, F.B., 1982. The mythical man-month: essays on software engineering. Addison-Wesley, Reading, MA.

Choo, C.W., Information management for the intelligent organization. Information Today/Learned Information, Medford, NJ, 1995.

Cohen, W.M., Levinthal, D.A., 1990. Absorptive capacity: a new perspective on learning and innovation. Administrative Science Quarterly (35), 128–152.

Cook, T.D., Campbell, D.T., 1979. Quasi-Experimentation: Design and Analysis Issues for Field Settings. Rand McNally, Chicago, IL.

Cook, T.D., Shadish, W.R., 1994. Social experiments: some developments over the past fifteen years. Annual Review of Psychology (45), 545–580.
Cooper, R.G., 1994. Third-generation new product processes. Journal of Product Innovation Management (11), 3–14.
Crowe, M., 1997. Intellectual capital for the perplexed. Harvard Management Update August, 3–5.
Davenport, T.H., Prusak, L., 1998. Working Knowledge: How Organizations Manage What They Know. Harvard Business School Press, Boston, MA.
Drucker, P.F., 1994. The age of social transformation. The Atlantic Monthly November, 53–80.
Drucker, P.F., 1988. The coming of the new organization. Harvard Business Review 66 (1), 45–53.
Finholt, T., Sproull, L., Kiesler, S., 1990. Communication and performance in ad hoc task groups. In: Galegher, J., Kraut, R. (Eds.), Intellectual Teamwork: Social and Technological Foundations of Cooperative Work. Lawrence Erlbaum, Hillsdale, NJ, pp. 291–326.
Galbraith, J.R., 1973. Designing Complex Organizations. Addison-Wesley, Reading, MA.
Galunic, D.C., Eisenhardt, K.M., 1994. Renewing the strategy-structure-performance paradigm. Research in Organizational Behavior (16), 215–255.
Grant, R.M., 1996. Towards a knowledge-based theory of the firm. Strategic Management Journal (17), 109–122.
Kiesler, S., Sproull, L., 1982. Managerial response to changing environments: perspectives on problem sensing from social cognition. Administrative Science Quarterly 27 (4), 548–570.
Lawler III, E.E., Mohrman, S.A., Ledford Jr, G.E., 1995. Creating High Performance Organizations: Practices and Results of Employee Involvement and TQM in Fortune 1000 Companies. Jossey-Bass, San Francisco.
Mankin, D., Cohen, S.G., Bickson, T.K., 1998. Teams and technology: extending the power of collaboration. In: Mohrman, S.A., Galbraith, J.R., Lawler III, E.E. (Eds.), Tomorrow's Organization: Crafting Winning Capabilities in a Dynamic World. Jossey-Bass, San Francisco.
Mintzberg, H., Raisinghani, D., Théorêt, A., 1976. The structure of unstructured decision processes. Administrative Science Quarterly (21), 246–275.
Nahapiet, J., Ghoshal, S., 1998. Social capital, intellectual capital, and the organizational advantage. Academy of Management Review (23)(2), 242–266.
Payne, J.W., 1982. Contingent decision behavior. Psychological Bulletin (92), 382–402.
Ruggles, R., 1998. The state of the notion: knowledge management in practice. California Management Review 40 (3), 80–89.
Simon, H.A., 1945. Administrative Behavior. Free Press, New York.
Simon, H.A., 1991. Bounded rationality and organizational learning. Organization Science (2), 125–134.
Simon, H.A., 1977. The new science of management decision. Prentice-Hall, Englewood Cliffs, NJ.
Steiner, I.D., 1972. Group process and productivity. Academic Press, New York.
Toffler, A., Toffler, H., 1995. Getting set for the coming millennium. The Futurist March–April, 14.
Tricker, R.I., 1992. The management of organizational knowledge. In: Galliers, R. (Ed.), Information Systems Research — Issues, Methods and Practical Guidelines. Blackwell Scientific Publications, Oxford, pp. 14–27.
Tsoukas, H., 1996. The firm as a distributed knowledge system: a constructionist approach. Strategic Management Journal (17), 11–25.
Tushman, M.L., Nadler, D.A., 1978. Information processing as an integrating concept in organizational design. In: Nadler, D.A., Tushman, M.L. (Eds.), Managerial Behaviour. 4th ed. Columbia University Press, New York, pp. 157–190.
Van de Ven, A.H., Delbecq, A.L., Koenig Jr, R., 1976. Determinants of co-ordination modes within organizations. American Psychological Review (41), 322–338.
Van Valen, L., 1973. A New evolutionary law. Evolutionary Theory (1), 1–30.
Yin, R.K., 1994. Case Study Research: Design Methods. 2nd ed. Sage, Thousand Oaks, CA.

Peter H. Gray has a background in electronic commerce and online information services and an MSc in Management Information Systems from Queen's University. He worked in information technology and management consulting positions before returning to academia, where he is a PhD candidate at Queen's University and a research associate with the Queen's Management Research Centre for Knowledge-Based Enterprises. He has published in the communications of the AIS, and has a sorthcoming publication in Decision Support Systems. His research interests include knowledge management and knowledge management systems, organization theory, information systems strategy, diffusion of innovation and social networks. He can be reached at pgray@business.queensu.ca.