Chapter 13

KNOWLEDGE MANAGEMENT IN ACCOUNTING AND PROFESSIONAL SERVICES

Daniel E. O'Leary
University of Southern California

INTRODUCTION

The purpose of this chapter is to (1) summarize and extend prior research in knowledge management; (2) overview the corresponding knowledge bases used in accounting, auditing, and taxation; and (3) point to new research opportunities. The frame of the manuscript is influenced in particular by its focus on knowledge management in professional service firm environments, such as the Big 5 accounting firms. In addition, this chapter highlights the role of alternative types of research in knowledge management.

Within the context of this chapter, knowledge management is defined as those efforts designed to (1) capture knowledge; (2) convert personal knowledge to group-available knowledge; (3) connect people to people, people to knowledge, knowledge to people, and knowledge to knowledge; and (4) measure that knowledge to facilitate management of resources and help understand its evolution. Capturing knowledge (effort 1) extends the traditional view of knowledge management to include emerging research interest related to when information or knowledge should be added to a knowledge management system (KMS) and how processes are reengineered to account for knowledge capture and access (e.g., O'Leary 1999a). Efforts 2 and 3 are derived from more traditional perspectives on knowledge management (see, for instance, O'Leary 1998c). Measurement (effort 4) is included as another dimension that is also emerging as an area of potential interest, particularly to accountants and related professionals because of the close tie to other traditional services they provide. These four types of efforts form the frame used in this chapter to examine knowledge management efforts and related opportunities for research.

Defining knowledge management has caused many heated discussions among academics and practicing professionals. Arguments are raised over "what is knowledge?" and "what is data?" As a result, no single definition of knowledge exists. According to Newell and Simon (1972, 810), the total knowledge available to a problem solver includes the knowledge state itself, access information, path information about how a given knowledge state was arrived at, access information to other knowledge states, and reference information. An important aspect of their definition is that knowledge includes information. This is consistent with usage in professional services, where KMS include so-called knowledge bases that are arguably simply information or data repositories. To avoid arbitrary classifications, Newell and Simon's (1972) definition is adopted within the context of this chapter as repositories of information and data that are classified as knowledge bases.

I acknowledge the extensive comments of two referees on an earlier version of this paper.

This chapter's review of the knowledge management literature examines the application of a range of methodologies. These methodologies include case studies, design research studies, and empirical examinations. Since knowledge management is relatively new, an important approach to eliciting and understanding issues is the case approach. In this setting, information is gathered about one or more companies or applications and is often attained through a series of interviews. Frequently, cases are used as the basis of generating an initial theory or explanation for certain events. Knowledge management also has roots in computer science and artificial intelligence. As a result, the design science methodology, typically implemented via classic computer-science-based design using a prototype or similar constructive approach, has been a commonly used methodology. Empirical research that analyzes data about KMS is another commonly used methodology. Studies that use this methodology are dependent upon systems being up and functioning so that historical data is available.

Before moving forward to the discussion in this chapter, it should be noted that the knowledge management literature is growing rapidly and in many directions. As a result, this chapter takes a narrow slice of the overall body of literature as it is most applicable to the accounting and professional services environment. There are a number of books available on the subject of knowledge management that interested readers may also wish to examine, including Davenport and Prusak (2000), Dixon (2000), and Krogh et al. (2000). Nonetheless, there has been only limited attention given to issues in knowledge management as they apply to accounting and professional services firms.

The purpose of this chapter is to overview the literature relative to knowledge management in accounting and the related professional services domain. The next section focuses on the knowledge bases that drive knowledge management and facilitates access to the knowledge. The following four sections focus on the four types of effort encapsulated by the definition put forth earlier—i.e., capture of knowledge, conversion, connection to existing knowledge, and measurement of knowledge. The chapter concludes with a brief summary and recognition of the limitations of scope necessitated by attempting to cover such a broad domain in a limited amount of space.

KNOWLEDGE BASES

Professional service firms employ a number of different types of knowledge bases as discussed in various case studies (Arthur Andersen 1997; Bartlett 1998; Chard and Sarvary 1997; Gladstone and Eccles 1991; Siboni 1997); however, little analysis of specific knowledge bases has been conducted. Accordingly, some of the primary issues, such as advantages and disadvantages are explored in this section. In the future, great benefit will come from researchers more fully expanding the issues associated with each type of knowledge base.

Knowledge bases include general information such as *news*, *internal expertise*, *who knows who*, *industry intelligence*, and *employees*. In addition, *frequently asked questions* and knowledge of *lessons learned* can also be helpful, in spite of their potential drawbacks. Among the Big 5 professional services firms, much of the corporate revenue process is driven by generating proposals for engagements in order to get the work, engagement reports, and supporting material after the proposal is accepted. As a result, knowledge bases that include *proposals and engagements* have been developed. Further, much of the Big 5 work is in the area of business processes, thus, they need to have knowledge and catalogs about *best practices* that can be employed for process improvement. Finally, within professional services firms there are functional divisions (audit, consulting, and tax). As a result, oftentimes knowledge bases are focused on specific *functional* areas. These knowledge base types are briefly described here along with particular issues of concern.

- *News Knowledge Bases:* There are a number of news knowledge bases (oftentimes termed *general knowledge bases*) that can be used to support firm activities. Typical related knowledge captured in KMS includes news feeds about world, national, local, and business news; internal

firm-developed news sources on corporate events and happenings; and articles on general management (e.g., *The Harvard Business Review*) and technical issues (e.g., *Information World*).

- *Who Knows Who Knowledge Bases*: Another key resource is knowledge of who in the firm knows whom outside the firm. This type of knowledge base was a key component of KPMG's original KMS design (Gladstone and Eccles 1991). For example, if someone in the firm knows the president of a potential client, that information can be very valuable as it may facilitate the sale of a product or service. There are a number of disadvantages to making this information generally available. First, since the knowledge is generally available, the person "who knows someone" may not know who is going to use the information or how they will use it. Second, except in rare circumstances (e.g., you know the president of a company), knowing who people should identify as being an important contact that should be included in the knowledge base is a difficult decision. Third, the people that you know may not want to be part of such a knowledge base and viewed as a knowledge asset. As a result, a number of disincentives also exist for people deciding whether to disclose whom they know.
- *Industry Intelligence Knowledge Bases*: KMS often include access to industry intelligence information (Siboni 1997). Competitor intelligence knowledge bases are likely to include information such as recent top management hires, recent client acquisitions and losses, or new service offerings. In addition, intelligence information may include a summary of competitor web pages. Unfortunately, the broad availability of information from sources such as the Web makes it difficult to keep pace with the information and changes in it. To accomplish this, one firm developed an approach to culling information from the Internet using intelligent agents (Steier et al. 1997). Their research would be categorized as design methodology, aimed at understanding how to build intelligent agents to perform industry intelligence analysis.
- *Knowledge About Employees:* Knowledge about employees is critical to all firms. Employee information may be kept in a structured group knowledge base or individually unstructured format (e.g., as a home page). There are advantages and disadvantages to both approaches. In the first case, employee issues of importance are identified *a priori*, and a manner in which the information is represented is established. Because it is semi-structured, the database is easily searchable and responsibility is established for maintaining the currency and accuracy of the database. In the second case, the individual user is likely to have more freedom as to what information is included on the page. However, the pages are likely to be nonhomogeneous in terms of the extent and frequency to which they are updated and in terms of the schemes used for information presentation and organization. This nonhomogeneous approach impacts the reliability and searchability of the information. There are at least three sources of such knowledge: (1) internal expertise knowledge bases, (2) human resources knowledge bases, and (3) employee home pages.

1. *Internal expertise knowledge bases* are designed to facilitate collaboration and to provide a repository for knowledge about who knows what. This allows individuals to draw on the expertise of others in the organization. Internal expertise knowledge bases attempt to funnel questions to the right people. For example, KPMG's early knowledge management efforts originated in what they called the "Who and Retrieval" System (Gladstone and Eccles 1991). Internal expertise knowledge bases have both advantages and disadvantages. The primary advantage is that people can be directed to the right person immediately. There are at least two disadvantages. First, focusing all the questions in one direction can be overwhelming to that person. Second, if all the questions are directed to a single source, expertise may not sufficiently grow throughout the firm.
2. *Human resource-based employee knowledge bases:* Typically, employee knowledge bases capture information about skills, experience, education, job level, etc. Employee knowledge

bases may be interfaced with job processing systems in order to determine current availability of someone who meets particular constraints. These knowledge bases are a part of almost all developed KMS.

3. *Employee Home Pages*: Employees may be presented with the opportunity to have their own home page. Since each individual does their own home page, even if there are standards, the aggregate information is unstructured.

- *Frequently Asked Question (FAQ) Knowledge Bases:* Frequently asked question (FAQ) knowledge bases are literally compilations of the questions most frequently asked about a particular subject. What are the advantages of FAQ knowledge bases? First, a question that is frequently asked is likely to be in the FAQ. Second, a substantial amount of knowledge is centered in a single location. What are some disadvantages of FAQ knowledge bases? First, knowing when a question belongs in one set of FAQ knowledge bases, rather than another is difficult. Second, FAQs will not include the unusual question that needs difficult-to-find information. Unfortunately, most users do not know when they have a question that is an FAQ or an SAQ (seldom asked question). Third, FAQs often are long lists, and actually finding the question that is being sought is extremely easier said than done. Some researchers have begun to explore solutions to these problems. FAQFINDER (Burke et al. 1997), developed to facilitate finding the right FAQ, is an example of design science research addressing these limitations.
- *Lessons Learned Knowledge Bases*: Lessons learned knowledge bases literally are designed to capture knowledge about when things go right (TGR), when things go wrong (TGW), and what the organization learned. For example, Ford has established such a knowledge base to provide a corporate memory (Stewart 1997). The Big 5 could (but have not at the time of this writing) use TGW and TGR to capture corporate knowledge about what went right or wrong on different engagements. What are the advantages of lessons learned knowledge bases? The primary advantage is that TGW and TGR can create an effective corporate memory. However, particularly with TGW, there may be more of a focus on responsibility. Thus, as noted by O'Leary (1998a, 56), "Generally TGR are easier to gather than TGW...since few employees are anxious to be associated with things that went wrong." Unfortunately, it may also not be clear when there was a TGR or a TGW. As a result, whether knowledge about an event will be captured or will be categorized correctly is not clear.
- *Proposal and Engagement Knowledge Bases*: The potential for proposal and engagement knowledge bases occurs on a spectrum, where at one extreme the knowledge base provides complete access to virtually all previous proposals and engagements. At the other end of the spectrum is the case where there are a few "sanitized" (e.g., name is removed) examples. Proposal and engagement knowledge bases provide a number of advantages (e.g., Gladstone and Eccles 1991). First, with access to the right proposal, writing another proposal for a related activity can be much easier. Second, a previous proposal can ease the process of costing a job and assist in determining necessary job steps for completion. There are also disadvantages of proposal and engagement knowledge bases. First, engagement knowledge bases make information that was done for a specific client available to other clients. Making results widely available can minimize payoff for the original client. As a result, clients often demand that no information about them be captured in the knowledge bases. Second, how can the user tell if the original proposal or engagement report was good? Unless the user is also an expert, s/he may not be able to determine appropriateness. Third, the more client-specific that the proposal and engagement materials are, the less the material is generally useable. Yet, a good proposal is likely to contain a good amount of client-specific information and utilize an approach that conforms to the client's specific needs. Fourth, users often cut and paste past documents—even accidentally including the previous client name in the new proposal or engagement report.

- *Best Practice Knowledge Bases:* Best practice knowledge bases provide users access to a wide range of so-called best practices or leading practices. Consider procurement as an example of a best practice. One approach to procurement information flow promulgated in Hammer (1992) uses only two documents, the purchase order and the receiving memorandum, dropping the invoice from the process. Procurement has been an active area for best practices (e.g., Cottrill 2000). Best practice knowledge bases may include diagrams outlining flows of information, war stories discussing the advantages and disadvantages of different best practices, lists of information about the particular best practice, and other resources. Best practice knowledge bases are complex and difficult to develop. Best practice knowledge bases are discussed in more detail in O'Leary (1998b, 2000b) and O'Leary and Selfridge (2000).
- *Functional Knowledge Bases:* Knowledge bases can also be functionally specific. For example, accounting has a number of specific knowledge bases to meet particular needs in accounting, e.g., financial reporting or tax. One of the first financial reporting knowledge bases was "Bean" (Business Executive's Accounting Network, which has evolved now as CFO Direct) provided by PricewaterhouseCoopers. Bean provided a clearinghouse for data, knowledge, and activities on accounting and auditing issues. It also provided related technology tools to facilitate use and searchable access of other knowledge sources (e.g., EdgarScan—Steier et al. 1997).

While a fairly extensive list of knowledge bases have been described in this section, various other authors have classified knowledge bases along similar and dissimilar lines. However, the classification provided here does describe the scope of contemporary knowledge base use in existing and emerging KMS.

From a research opportunities perspective, few of the issues associated with different types of knowledge bases have been fully flushed out. The discussion in this section lays out a number of advantages and disadvantages of most such knowledge bases, but additional research is still needed—particularly research addressing the disadvantages and possible solutions. For example, what happens when a firm introduces a given type of knowledge base? Redundancy issues also occur in knowledge bases—particularly with employee expertise information—as data and information frequently appear in multiple knowledge bases: What settings are likely to lead to redundancy? What are its costs and benefits? How can the redundancies be reduced?

Client privacy issues also drive some major concerns with capturing knowledge in KMS. An emerging concern is what belongs to the client and what belongs to the professional services firm. Should proposals and engagements be made generally available to staff within the firm or should each engagement be handled independently?

Different functional knowledge bases were also briefly discussed above. Functional knowledge bases evolve over time. Comparing different versions can possibly provide insight into change forces and how knowledge or knowledge about knowledge bases evolves over time. Another important set of issues is how to manage that change.

CAPTURING KNOWLEDGE

Knowledge capture is the first step in knowledge management. In order for knowledge to be captured, a set of mechanisms must be in place to determine when new knowledge should be captured. Further, in order to facilitate knowledge capture, the process of knowledge management needs to be embedded in specific processes. As a result, reengineering efforts need to consider knowledge management and conversely.

The concept of triggers is used to determine when to capture information and knowledge. In transaction-processing systems, a well-defined set of events is used to determine when new information should be added to the database. In KMS, triggers are also needed so that new knowledge is added. The nature of the trigger varies based on the specific knowledge base for which the trigger is being

developed. For example, in the case of an *employee knowledge base*, if a new employee joins the firm, then new information should be added to the knowledge base. Similarly, whenever an existing employee works on a different activity or job, or whenever the employee gets new training, additions should be made to the knowledge base.

Triggers for some knowledge bases are not as easily generated. For example, in the case of *lessons learned knowledge bases*, knowing when to gather new knowledge is not always clear. In addition, whether knowledge about an event will be captured and categorized correctly is not clear. Accordingly, a system needs to provide triggers to capture events and categorize whether an event was a TGR or a TGW.

To illustrate, one large Fortune 500 company employs traditional management accounting approaches as a means of triggering the event and categorizing it. The budget is used to trigger when an event occurs. A category that exceeds budget is categorized as a TGW, while a category that is under budget is categorized as a TGR. In this setting the relationship to the budget triggers the event of gathering an explanation that would be captured and put in the system. This approach does have both advantages and disadvantages. As a result of the automatic trigger, when an item should be captured for the *lessons learned knowledge base* is certain. Similarly, the conditions for classification as TGW vs. TGR are predefined and clear-cut. However, the clear-cut nature of the definition may actually end up defining some situations as TGW when they are really TGR. For example, research and development expenditures may exceed budget, because just a small percent more were needed to fully flush out a product. As a result, over-budget expenditures may yield a successful product development that would be categorized as a TGW. This same approach can be extended from a reliance on the budget to include other management accounting approaches. For example, traditional variance analysis provides a similar approach to triggering capture of events and how to categorize them.

Another challenge of effective knowledge management is the necessity that the KMS be embedded in the processes that the KMS is designed to support. Perhaps the best example of this is customer support where it is not unusual for the process to automatically capture information related to which client called, who answered the call, the basic nature of the call, the ability to resolve the caller's problem, etc. As a result, a substantial amount of information and knowledge is available to future users who need to know about the client, the product they purchased, the kind of questions the client asked, etc.

Unfortunately, anticipating how knowledge management can be readily embedded in many audit, tax, or consulting processes is more difficult. It is not just a matter of embedding the KMS. O'Leary (1999a) addresses this issue and presents a model of the stages of interaction between reengineering and knowledge management (see Figure 13-1). He argued that both process reengineering and knowledge management must be accounted for in the process design, illustrating the concepts with a case study. Thus, in terms of Figure 13-1, design for embedded KMS would push a system from the lower left hand corner to the upper right hand corner.

There is again a broad range of research opportunities related to capturing knowledge for KMS. One of the most important issues is how can knowledge management be embedded in particular activities, whether they are audit, tax, or consulting in nature. Oftentimes knowledge management and reengineering of processes must be considered simultaneously. The discussion in this section also raises the issue of triggers, a critical but overlooked aspect of knowledge management. Research needs to address how firms effectively can and do trigger capture of knowledge for the KMS. As noted above, management accounting would seemingly be a particularly effective source of metrics for triggers.

FIGURE 13-1
Stages of Interaction between
Reengineering and Knowledge Management

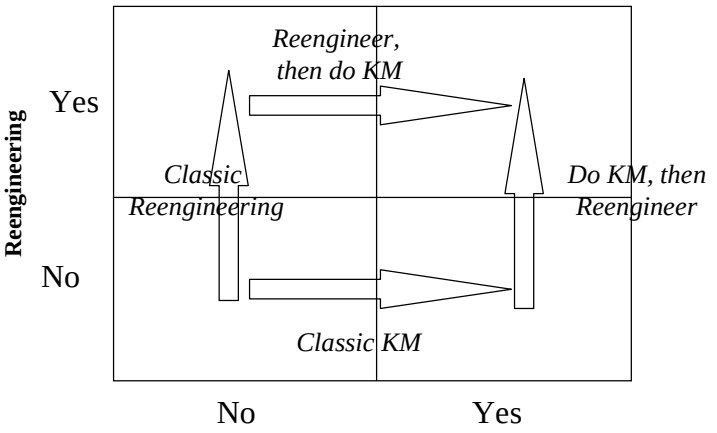

CONVERSION

As discussed in O'Leary (1998c) conversion uses knowledge harvesting and sharing to convert personal knowledge to group-available knowledge and conversion takes text and data and uses knowledge discovery to generate knowledge. Knowledge sharing is one of the most important aspects of knowledge management; however, it has received only limited attention. Further, most of the attention to date has been focused on specific cases (for example, see discussion of the World Bank in ebusinessforum [2000]). Orlikowski (1993) provided some of the first case-based analyses associated with the problems of knowledge management in a professional services firm in an examination of the use of Lotus Notes. As part of that investigation, she identified barriers to conversion of knowledge. For example, as noted in one of the interviews that she gathered:

> Power in this firm is your client base and technical ability.…It is definitely a function of consulting firms. Now if you put all this information in a notes database, you lose power. There will be nothing that's privy to you, so you will lose power. It's important that I am selling something that no one else has. When I hear people talk about the importance of sharing expertise in the firm, I say, "Reality is a nice construct." (Orlikowski 1993, 246)

Understanding what forces may lead to a sharing of knowledge and what forces may lead to a lack of sharing is important. These forces will in turn facilitate or hinder conversion of personal knowledge to group available knowledge. Unfortunately, research beyond Orlikowski (1993) has been limited; more research is needed in this area.

Conversion issues provide a great deal of opportunity for researchers. For example, how do professional service firms facilitate knowledge sharing? What examples can be identified in professional service firms' practices that illustrate such knowledge sharing? How do the more effective or efficient ways of facilitating sharing actually work in practice? Are there other methods that have yet to be identified that could enhance such knowledge sharing?

There is also a limited body of design science research in the area of conversion of text and data to knowledge. For example, Steier et al. (1997) developed a system called ODIE (On Demand Information Extractor). ODIE reads over a thousand news stories nightly to generate knowledge about management changes. Professional service firms can then use that knowledge to determine

which firms to target as potential clients. The approach that they use can be generalized to other applications where they can employ knowledge about specific events, syntactical patterns, and stylized articles. For example, Huffman and Steier (1995) investigated management changes, major shareholder changes, and audit committee changes as other sources of such knowledge.

In addition, there has been limited research on how to convert data to knowledge to meet the needs of professional service firms. Such research could use design science methods by exploring alternative machine-learning approaches. In addition, research could focus on the use of such generated knowledge, including the extent to which new clients were added as a result of availability of knowledge.

CONNECTION TO EXISTING KNOWLEDGE

In order to connect people to knowledge, knowledge needs to be placed in a usable format. Generally this entails either developing computer-based systems with hyperlinked data or alternative organization structures based on a common language to facilitate use. Connection also includes the determination of how to deploy the knowledge.

Most technologies for knowledge representation in a KMS environment are based on Internet browsing platforms. Such an organization usually employs a series of hypertext links that link data and information from one web page to another. For a KMS to be successful, the pre-defined hypertext links must be usable and provide an appropriate set of links that allow the user to get to the embedded knowledge that is required for a given decision. There are two strategies to development of such links. The first strategy is to include only a few selected links to ensure that the browser keeps the user on track for solving the problem. The second strategy is to enable all possibly related links on a given page so that the user can sift through and determine which information is key to the decision at hand. The disadvantages to the two strategies mirror the other's advantage. The first strategy may result in no useful links, while following the second strategy, the user could get further and further off course from the original objective. Little empirical research has addressed these two competing strategies (or some in-between strategy), particularly in a professional services setting.

Knowledge can also be organized and represented for not only human-based consumption, but also computer-based consumption—an area of particular interest in the domain of artificial intelligence (AI). One area of knowledge management that has found substantial application in AI is the area of *best practices*. There are a number of approaches for capturing knowledge to facilitate reengineering. O'Leary (2000a) modeled *best practice* knowledge using Bayes' nets, uncertainty factors, rule-based systems, and case-based reasoning approaches. In particular, with case-based reasoning, *best practices* are represented as cases with a number of dimensions ("slots") used to capture different characteristics of the *best practice*. For example, in a prototype discussed in O'Leary (2000a), applicable characteristics included industry, company, process affected (e.g., accounts payable), resources affected (e.g., inventory), technologies used (e.g., scanning or bar codes), side of supply chain affected, and other factors. Finding a *best practice* that meets requirements was then a function of matching firm needs to the existing case base of *best practices*. Although each approach was successfully used to model *best practices*, in a comparison of those methodologies, O'Leary (2000a) concluded that case-based reasoning had a number of advantages over other approaches. The case-based approach provides a representation that is more detailed and considers a broader range of factors. Cases can be easily added as new instances of best practices become available. Further, cases can be readily generated from existing knowledge bases of *best practices*.

In order to facilitate knowledge connection, be it human-based or computer-based usage, KMS require a common language—also referred to as ontology. Gruber (1993) defines ontology as an explicit specification of a conceptualization. Ontologies define knowledge components with sufficient specificity that users can find the information that they need, whether through direct search or search engine. As such, ontologies provide multiple agents, whether human- or computer-based, with the ability to communicate without ambiguity.

One area where there has been work in the area of ontologies for knowledge management is that of *best practices*. In particular, probably using the work of Porter (1985), at least two firms have begun to establish what they refer to as "common languages." Those firms have made representations of their *best practice* knowledge bases available. General categories of knowledge are given for Price Waterhouse (Figure 13-2) and Arthur Andersen (Figure 13-3). The two categories are somewhat similar. For example, Price Waterhouse has *best practices* in each of the categories summarized in the figure, such as "Perform Marketing and Sales." Although both seem to have a common heritage in Porter's (1985) value chain, empirical research finds a number of differences between the alternative ontologies (O'Leary 2000b). In addition, analytic research, based on Arrow's well-known impossibility theorem, finds that there is no best practices ontology that is optimal for all firms (O'Leary 2000b).

A related issue is the question of how to deploy knowledge bases once they are in place. There are a number of advantages to deploying KMS through an intranet. First, an intranet provides wide availability. Second, intranet access can be restricted through control procedures that provide general security of the knowledge in the KMS. However, the intranet also makes the knowledge bases accessible across the firm by users from multiple lines of business, e.g., consulting, audit, and tax. To the extent that there is cannibalism across product lines or by one office of another office's clients, limited diffusion may be desirable. Perhaps because of these concerns or others, professional service firms continue to make extensive use of CDs. With a CD, the knowledge does not have to be generally available across the entire firm. Rather, CDs can be used to limit knowledge diffusion through limited distribution.

Given the importance of knowledge deployment to an overall connection strategy, an emerging issue that needs exploration by researchers is how to most effectively deploy KMS to the right people, while controlling access by other groups. With all the schisms in professional services firms, limiting deployment can be as big a problem as making sure that the right people have access. There has been limited prior case-based research on this issue, and virtually no empirical research.

MEASUREMENT OF KNOWLEDGE

FIGURE 13-2
Price Waterhouse Best Practices Knowledge Base

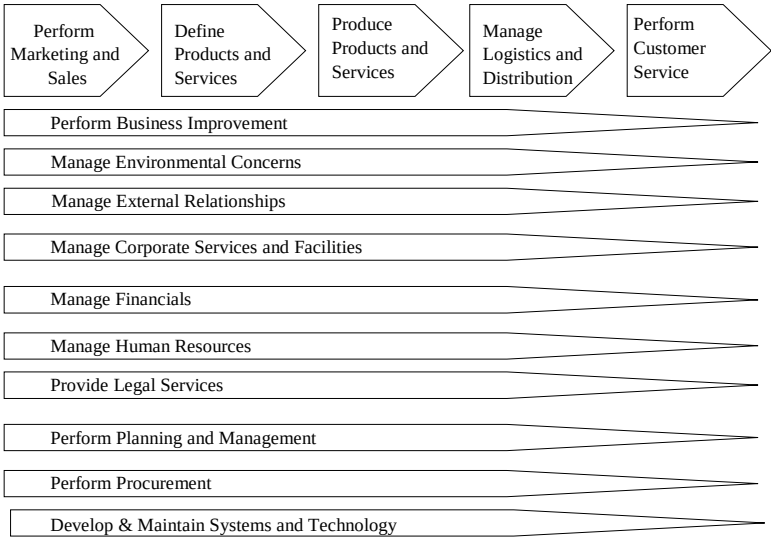

KnowledgeView Multi-Industry Process Technology (Price Waterhouse)

FIGURE 13-3
Arthur Andersen's Best Practice Knowledge Base

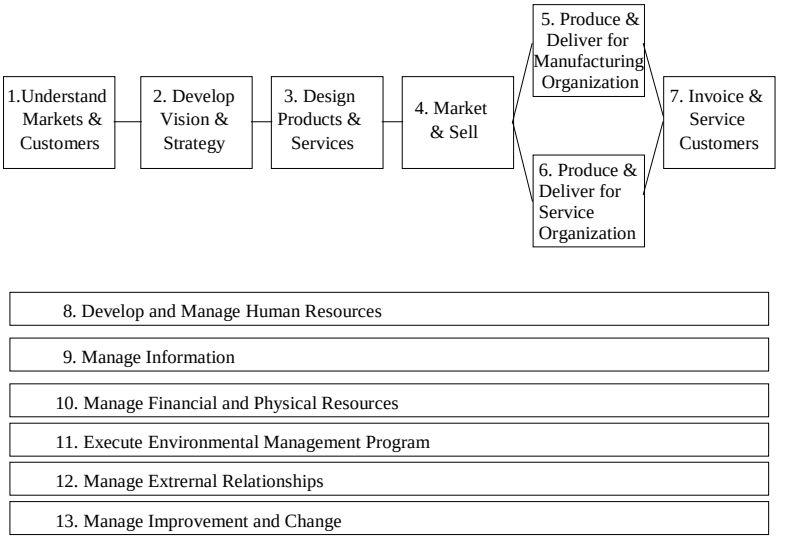

Global Best Practice Classification Scheme (Arthur Andersen)

Ultimately, the value-added from knowledge management can only be known if effective knowledge measurement processes are in place to facilitate control, value assessment, and cost-benefit analysis. For example, Davenport et al. (1998) proposed that knowledge management projects are successful, only when appropriate business indicators provide "some evidence of financial return either for the knowledge management activity itself…or for the larger organization."

As a result, a number of measurement approaches have been proposed. For example, Bassi (1999) identifies 12 ways to measure intellectual capital including, for instance, balanced scorecard, subsystem performance (e.g., Dow Chemical measured an increase in licensing revenues from control of patents), and benchmarking. Unfortunately, at this point limited empirical analysis investigates which companies use which approaches, or the effectiveness of alternative approaches.

Which measurements are used appears to depend primarily on what the user or management wishes to measure (O'Leary 1999a). For example, management may be concerned with better understanding about how the KMS is facilitating collaboration. In that case, measures such as percentage of time spent collaborating and percentage of groups actually using collaboration systems can be used. If the issue is the extent and connectedness of the knowledge, then measures such as number of links connecting knowledge and complexity of network connecting knowledge in different areas can be used. In this latter case, analytic research, such as graph theory, can be used to study performance.

Given the expertise that accountants have in measurement issues, this may be a particular good area of opportunity for AIS researchers. Very little specific knowledge management research has been completed to date, yet the ability to measure the degree of success is critical to continued investment in KMS and in determining approaches for enhancing the value of KMS.

SUMMARY

This chapter has reviewed the use of knowledge management in accounting with a focus on how professional service firms can and do manage their knowledge. Within the context of this discussion, knowledge management was defined as the methods used to facilitate capture of new knowledge; conversion of data and personal knowledge to group available knowledge; connection of people to people, people to knowledge, knowledge to people, and knowledge to knowledge; and measurement of knowledge.

The first part of the chapter focused on the different types of knowledge bases for KMS used in professional services firms. For these different knowledge bases, the advantages and disadvantages were discussed as precursors to the identification of potential research opportunities. Subsequently, the chapter shifts its focus to the component efforts identified in the adopted definition of knowledge management. Capturing new knowledge established concerns regarding the role of triggers in knowledge management and a discussion of issues that relate to how to capture knowledge from processes. Converting personal knowledge to group knowledge faces limitations that might be overcome through methods identified by future researchers. In addition, converting knowledge using machine-learning approaches seems to have received only limited attention by the professional services firms and is another area needing additional examination. Connection of knowledge can take a number of approaches with perhaps the most interesting to researchers being the use of AI to connect users to knowledge. Finally, some selected measurement issues in knowledge management were reviewed.

Although the focus here is on these four effort dimensions of knowledge management, other critical knowledge management issues exist—but are beyond the scope of this chapter. These other issues include for instance collaboration (Orlikowski 1993) and creation of new knowledge (Bartlett 1998). In addition, knowledge management organization, organization culture, performance incentives, and organizational structure issues fall into the organizational issues domain and may be fruitful areas of research. These organizational issues are outside the scope of this paper, but can be critically important to successful knowledge management (see for instance, Bank 1996; Chard and Sarvary 1997; March and Garvin 1997). The interested reader wishing to pursue research in knowledge management may also wish to explore these other areas in greater detail.