EXPLORING THE ROLE OF NETWORK EFFECTS IN IT IMPLEMENTATION: THE CASE OF KNOWLEDGE MANAGEMENT SYSTEMS

Jürgen P. Bansler, Erling Havn
Center for Tele-Information
Technical University of Denmark, Building 371
DK-2800 Kongens Lyngby
Denmark
Phone: +45 4587 1577
Fax: +45 4596 3171
bansler@cti.dtu.dk, havn@cti.dtu.dk

ABSTRACT

One of the key themes in knowledge management is the role of information systems designed to facilitate the sharing and reuse of knowledge – often referred to as knowledge management systems (KMS). However, the intended users of such systems are often reluctant to use them, and implementation failures are common. While prior research offers important insights into the problems and practices of IS implementation in general, little is known about the special problems involved in implementing KMS. This study seeks to better understand the process of KMS implementation and establish a theoretical framework for examining the underlying dynamics of adoption and use. We suggest that efforts to implement KMS need to be sensitive to the social interactions and the collective sensemaking of the intended users if they are to be successful. In making this argument, we build on recent research within the field of network economics and highlight the concept of network effects*. Using an exploratory case study as an illustration, we argue that KMS exhibit strong network effects and that these create positive feedback loops that complicate the implementation process. Our conclusion is that the concept of network effects offers an interesting and useful analytic perspective for understanding the implementation of KMS in organizations. Implications of using this theoretical lens for both research and practice are discussed.*

1. INTRODUCTION

There is general recognition among both researchers and practitioners that knowledge has become an important organizational asset and a basis for sustainable competitive advantage. Many organizations are developing IT systems designed specifically to facilitate the flow of knowledge among individuals and groups across organizational and geopgraphical boundaries (Markus 2001). Such systems are often referred to as organizational memory systems or knowledge management systems (KMS) (Alavi & Leidner 1999). KMS are expected to improve organizations' productivity, flexibility and innovative capability by enabling organizational members to share, integrate and reuse knowledge more effectively. In practice, however, results have been mixed. Some knowledge sharing systems have been successful (Brown & Duguid 2000), but implementation failures are common and the intended users are frequently reluctant to use the systems (see e.g. Orlikowski 1992).

A large and growing body of IS research has addressed the complex issue of IT implementation in organizations (Lai & Mahapatra 1997). Many studies deal with management of IT implementation and factors critical to the success and/or failure of implementation projects. While this research has shown that implementation failure may be caused by many different factors – including bad systems design, organizational culture, and political processes – it does not address the special challenges involved in implementing KMS.

Implementation of KMS requires very different approaches than are necessary in other IT implementations. There are two reasons for this. First, successful implementation of an organizational knowledge sharing system depends on the willingnes of users to participate actively in the process. In fact, the very label of "user" is somewhat inappropriate in the context of knowledge sharing systems, as users are both contributors and beneficiaries of the system (Alavi & Leidner 1999). Secondly, users of a knowledge sharing system are mutually dependent and the benefits experienced by one user are contingent on the number and behavior of other users. For instance, a shared knowledge repository to which few people contribute their knowledge is without value – but when many people choose to share their knowledge, the repository may become a highly valuable resource (Brown & Duguid 2000). As a consequence, the successful implementation of a KMS requires the active participation by a "critical mass" of users almost from the start of the implementation. By examining this interdependence among users more closely, we can gain much insight into the difficulties associated with the implementation of systems for knowledge sharing.

The concept of interdependence and the phenomenon that the benefits of a system (or network) increase with the total number of users are well-known in telecommunication markets and markets for durable consumer goods. Economists refer to this as "network externalities" or simply "network effects." Network effects have important implications for market dynamics: Technologies or systems subject to strong network effects tend to get off to a slow start and, then, either reach "critical mass" and exhibit explosive growth, or if they fail to do so, disappear (Shapiro & Varian 1999). This pattern results from positive feedback: as the customer base grows, the benefits increase and attract more and more users (and vice versa, if the customer base starts to shrink, then benefits will decrease and users defect).

While economists have been studying network effects and their implications for market dynamics for decades, IS researchers have paid little attention to the role of network effects in organizations. About ten years ago, a few papers developed the notion of "critical mass" as an approach to understanding the adoption of interactive communication technologies, like email and voice mail (Markus 1987, 1990, Rice 1990, Rogers 1990). Since the publication of these papers, scholarly interest in network effects and their impact on the implementation of new technologies in organizations appears to have dwindled and we are not aware of any attempts to analyze the implementation of KMS from this perspective.

In this paper we draw on recent developments within network economics to analyze the role of network effects in relation to the implementation and use of systems for knowledge sharing in organizations. We focus on the sources of network effects and their influence on implementation dynamics and adoption patterns. In order to explore the concept of network effects in more detail and provide a concrete illustration of its application and value, we present some findings from an empirical study of the implementation of a shared "best practice" repository in a large organization. We conclude with a discussion of the implications of our analysis for IS research and practice.

2. THE CONCEPT OF NETWORK EFFECTS

The classical example from economics of network effects is the telephone (Rohlfs 1974). The telephone provides no benefits to isolated users, but when there are many other users, the benefits increase significantly (Economides & Himmelberg 1995; Shapiro & Varian 1999). The network effects give rise to direct positive feedback: as new subscribers join the telephone network, the incremental utility of the service increases, thus encouraging marginal nonadopters to adopt, causing further growth in the telephone system. (See figure 1a.)

Network effects are, however, not confined to communication networks. Many other products and services display network effects, though usually in less-obvious ways[1] (Economides 1996, Katz & Shapiro 1985). Consider, for example, the owner of a personal computer. At present there are two competing technologies, Macintosh and Wintel. These systems are incompatible with one another; programs developed for the Macintosh platform cannot run on the Wintel platform – and vice versa. Because of the increasing returns to scale in the provision of software, the owner of a personal computer will find a greater variety of applications for his machine if more computers are sold using his technology. He also benefits from the ability to exchange programs and files with other users of compatible machines and from superior service that may be available for the computer technology with the larger installed base of machines (Katz & Shapiro 1986). In this case, the chain of cause-effect relationships is longer and the positive feedback loops more indirect, but not less significant. (See figure 1b.)

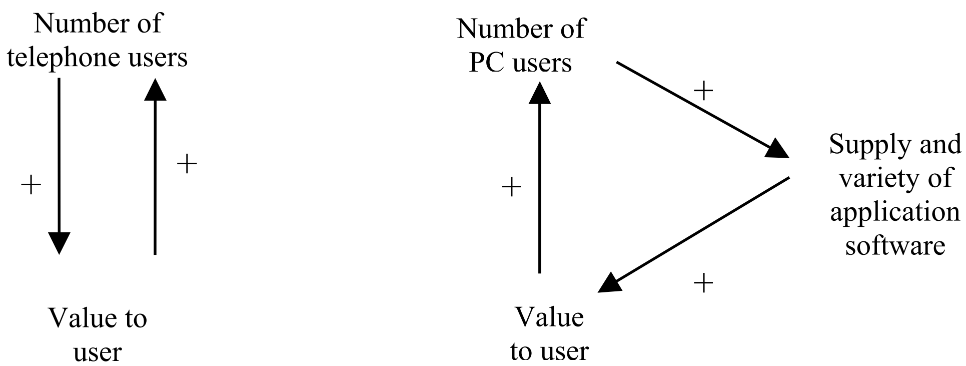

Figure 1. Examples of (a) a direct positive feedback loop (telephone network) and (b) an indirect positive feedback loop (the network of PC users)

The notion of positive feedback is crucial to understanding the adoption of new technologies in markets where network effects are significant. The typical pattern involves an S-shaped curve with three phases: (1) flat during early introduction, then (2) explosive growth during "takeoff," followed by (3) leveling off as saturation is reached. (Rogers 1990, Shapiro & Varian 1999) See figure 2. In the beginning, when the number of users is moderate, the actual use value is limited and only users with special interests or needs will adopt the innovation. But as the number of users grows, more and more users find adoption worthwhile. Eventually the number of users achieves a certain point where

[1] Katz & Shapiro (1985) mention three important sources of network effects: (1) The network effect may be generated through a direct physical effect of the number of purchasers on the product's quality. Telephone networks are the prototypical example. (2) Network effects may also occur indirectly. For example a consumer purchasing a personal computer will be concerned with the number of other consumers purchasing similar hardware because the amount and variety of software being supplied for a particular computer will be an increasing function of the number of hardware units that have been sold. They call this the "hardware-software paradigm." (3) Network effects also arise in relation to durable goods, e.g. an automobile. In this instance, the quality and availability of postpurchase service will depend on the experience and size of the service network, which may in turn vary with the number of units of the good that have been sold. In addition to these major sources of network effects they also mention three more subtle ones. These include: (4) the fact that product information is more easily available for more popular brands; (5) the role of market share as a signal of product quality; and (iii) purely psychological, bandwagon effects.

adoption begins to accelerate and explosive growth sets in. Inspired by physics[2], this point is usually called the *critical mass* point and the key challenge is to reach this point. Once the installed base is large enough, the market will build itself.

However, many new technologies fail to obtain critical mass and simply flop. The problem they face is sometimes referred to as the "chicken and the egg" paradox: many users are not interested in adopting the technology because the installed base is too small, and the installed base is too small because an insufficiently small number of users have adopted the technology.

In both cases, user expectations are critical. The size of the installed base (i.e. the total number of users) of the technology is a key factor in consumer decisions about whether or not to adopt a new technology, but *expectations* as to how widespread the technology will become in the future are equally – if not more – important (Economides 1996, Shapiro &Varian 1999). As a result, expectations about success or failure tend to become self-fulfilling prophecies.

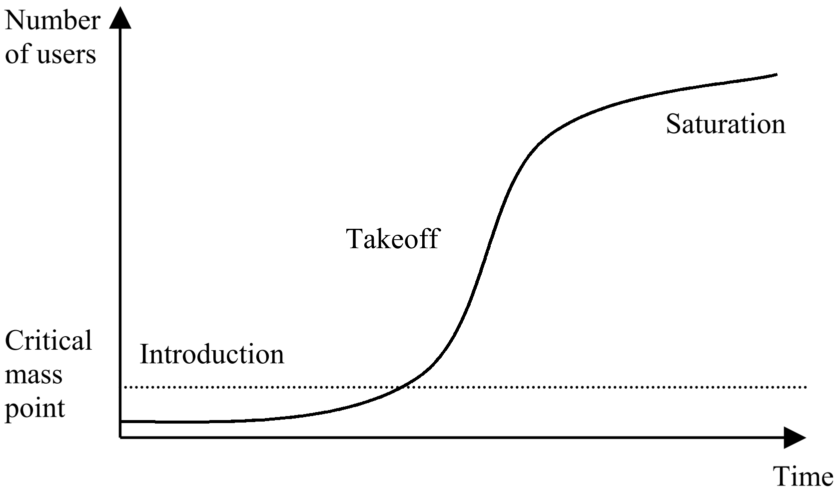

Figure 2. Adoption dynamics in network markets (successful)

3. NETWORK EFFECTS IN ORGANIZATIONS

Although the concept of network effects usually is associated with the study of the structure and dynamics of markets for telecommunication services and durable consumer goods, it is also applicable in an organizational context. At a time when organizations are becoming increasingly "wired," with large scale implementation of intranets and growing importance of organizational memory and knowledge management systems, the pervasiveness and impact of direct and indirect network effects grow.

Although we do believe that network effects often play a significant (and growing) role in organizations – just as they do in the market – there are, of course, differences between introducing an innovation in the marketplace and implementing a new technology in an organization. Individual consumers are – at least in principle – free to choose what to buy and when to buy it, whereas the

[2] According to Rogers (1990), the concept of critical mass originated in physics, where it was defined as the amount of radioactive material necessary to produce a nuclear explosion. An atomic pile "goes critical" when a chain reaction of nuclear fission becomes self-sustaining. If the amount of fissionable material falls below the critical mass the reaction will peter out.

employees in an organization have less autonomy and discretion because their managers have the authority to mandate the use of specific systems and tools. Thus, in principle, organizations can (and sometimes do) implement new systems and tools simply by mandating use and in this way eliminate the influence of network effects.

In practice, however, implementing KMS by “brute force” is a poor strategy. First, mandating usage and coercing users is likely to create hostility and tension, and may stimulate resistance against the technology. Second, successful implementation of KMS requires that the intended users are motivated to share their knowledge with others and participate actively in the implementation process. In most cases, implementation success or failure will therefore depend on the perceived benefits and costs to individual users (now and in the near future)

In order to explore the concept of network effects in more detail and provide a concrete illustration of its application, we now present some findings from an empirical study of the introduction and use of a best practice database in a multinational biotech firm.

4. CASE STUDY

4.1 Research Site and Methodology

The field study investigated the introduction and use of a best practice database in a large, biotech firm, Beta Corporation (a pseudonym). Beta Corporation develops and manufactures a range of medicinal drugs and industrial chemicals. With production facilities, research centers, and sales offices in more than 60 countries, Beta employs more than 14,000 employees. Beta is a knowledge-intensive company with strong ties to universities and research hospitals.

In 1996, top management at Beta decided to implement a corporate-wide best practice database, SHARE (a pseudonym), as part of a strategy to accelerate and broaden the sharing of managerial knowledge across functions and organizational units. The objective was to create an electronic repository that could enable Beta’s managers to exchange “best practices” over the company’s newly established intranet. The idea was that managers who had identified an effective way to perform a process would submit a description of the practice to the common on-line repository so that other managers could quickly learn about it. The goal was not to specify all the needed knowledge about a best practice but rather to provide enough information to allow an interested manager to evaluate the practice and contact the author for more details.

The research described here focused on the implementation and use of the best practice database over a three-year period, from its introduction in January 1997 until the end of 1999. We investigated how the best practice database was designed and redesigned several times during this period, how it was implemented in the organization, and how managers reacted to it.

A qualitative and exploratory case study approach was used to collect and analyze the data (Starke, 2000). Detailed data collection was conducted through interviews and document analysis. In addition we had access to the SHARE system and its content. Most of the interview data came from a series of in-depth and unstructured discussions with the project manager who was responsible for the design and implementation of the SHARE system. These were supplemented by ten more structured interviews with system designers and managers (as users of the system). Each interview lasted between one and two hours and all interviews were tape-recorded and transcribed.

4.2 Course of events

Senior management at Beta wanted to identify and foster best practices among their managers – not by imposing procedures from above, but by responding to the inventive improvisational ways people actually get things done. The basic idea was to foster invention and creativity by allowing managers to improvise and experiment with new ways of working, and then identify and circulate the best ideas and practices corporate wide. The SHARE project set out to create a "best practice database" to store useful ideas and spread them around the world. The introduction and use of the SHARE database can be divided into three distinct phases:

Phase I – Initial development and introduction of the system. The first version of the database was developed in-house and was ready for "take-off" in January 1997. The project group, backed by top management, launched a major campaign to inform managers at all levels about the system and encourage them to submit suggestions. Through a series of meetings, articles in the company newsletter and "banner ads" on the intranet, the project conveyed the message that this was a high profile initiative and that top management wanted everybody to participate actively in the creation and use of the database.

Top management emphasized that their goal was to create a high quality database. To ensure this, a rigorous, centralized review process was established. Suggestions, so-called best practice "candidates", were submitted to a committee of 14 experienced management consultants (i.e. internal consultants, usually former managers or senior managers who had been with the company for many years) who then evaluated the submissions.

However, it was almost immediately clear that this review procedure was problematic: The evaluation committee turned out to have very high standards and during the first round of reviews it rejected more than half of the submitted proposals. When managers learned about this, they were shocked and simply ceased to submit new proposals. As the SHARE project manager said later, "the first time this committee reviewed the submissions, it rejected 8 out 10 proposals – and that is something that kills people's motivation [to submit proposals for best practices]."

Phase II – Elimination of the review process. After three months, top management decided to drop the review process altogether in an attempt to revive the database and regain support from Beta's managers. After the review process was abandoned, managers could put a new BP in the database, without asking permission or any formalities. The concept of Best Practice was also substituted by the less ambitious concept of "Better Practice".

In the months that followed, corporate headquarters launched a new campaign to promote the system. "We had to get on our hands and knees to persuade them [Beta's managers] to participate," the project manager recalled. After a slow start managers hesitantly began to contribute to the database, and from April 1997 to October 1998 they put 368 new "Better Practices" in the database. The number of submissions per month peaked in the middle of this period (in December 1997 with 36 new BPs registered in the database) and then gradually decreased (in October 1998 the number had dropped to 9 submissions per month.)

Thus, after eliminating the formal review process, Beta succeeded in revitalizing the system and getting managers to build the database. However, few managers showed any interest in searching the database. During 1998 it was clear to the project manager that virtually no one used the database. When asked why they did not use the database, many managers explained that the design of the user interface made it difficult and time-consuming to search the database, and that, in general, the quality of the BPs in the database was poor. The lack of users (and the fact that the number of new submissions had begun to fall again after December 1997) made it clear that the project was in a critical condition and that something had to be done to save the project from failure.

Phase III – Redesign and reintroduction of a review process. As a response to the crisis, corporate management made two important decisions. First, they allocated substantial funds to the development of a new, improved user interface to the database. The new version of the software was developed by Beta in cooperation with a software house specializing in the design of Web-based user interfaces. The new design was more elegant and visually attractive and a series of usability tests showed that it was far easier to learn and use than the previous interface.

Secondly, corporate management decided to reintroduce a formal review process to improve the quality and usefulness of the BPs in the database. In ligth of Beta's previous experience, this time they deliberately decentralized the review process in order to avoid the same negative reactions as with the former evaluation committee. Instead of the central committee, corporate headquarters appointed 30 so-called "moderators" – leading specialist within their field of expertise – from different parts of the organization to review incoming submissions. Each moderator would only review submissions within his or her own fields of expertise. The idea was to establish a *peer* review process similar to the ones associated with the scientific community. By appointing reviewers who were closer to the "end users" of the database and by stressing the principle of peer review, top management hoped that people would be less intimidated by the review process and more willing to submit proposals for new best practices. The leader of the project put it this way:

> We have decentralized the ownership [of the database]. And I think that is going to make a big difference, because the people who are going to "clean up" now are no longer some distant committee with police caps, but some of their own people. Therefore, they have to accept it and also to participate actively [by submitting new proposals].

The new version of the system and the new review process were introduced in November 1998 and, once again, corporate headquarters launched a big campaign to promote use of the database. The result was, however, disappointing. Although virtually everybody agreed that the new version of the system was easier to use and that the new review process was a good idea, people still did not use the database. To make things worse, the flow of input, in terms of new submissions, slowly ebbed by spring 1999.

In June, the project manager arranged a meeting with the moderators to evaluate the situation. The meeting unanimously concluded that "the SHARE document base does not contribute significantly to sharing of knowledge" in Beta, and a few months later, top management finally – after three years – decided to close down the database.

5. DISCUSSION: REVISITING THE CASE IN LIGHT OF NETWORK EFFECTS

What went wrong here? We suggest that network effects and positive feedback loops, unrecognized by the main actors themselves, can account for much of what happened in this story. In particular, we believe that the contraproductive way in which the issues of quality and review procedures were handled contributed substantially to the failure of the system.

The SHARE database is an example of a knowledge sharing system with strong, indirect network effects. The users of the system can play two different roles, either as *contributors* of new best practices or as *information seekers* searching the database for useful ideas and processes. Of course, the same person can (and ideally should) play both roles – sometimes submitting new proposals and at other times seeking new insights and inspiration from the database. As we will show in the following, there is a significant interdependence between the behavior of contributors and information seekers.

5.1 Sources of network effects

The value of the SHARE system to information seekers obviously depends on the number of BPs in the database. If there are only few BPs in the database, its value is very limited, but when more BPs are added, the chance of finding something of interest increases (depending, of course, on the quality of the BPs, but we will postpone this discussion for a moment.) The number of BPs in the repository, in turn, depends on how willing managers in Beta are to take on the role of contributor and submit new proposals. The benefits of the database to information seekers are thus contingent on the number of contributors.

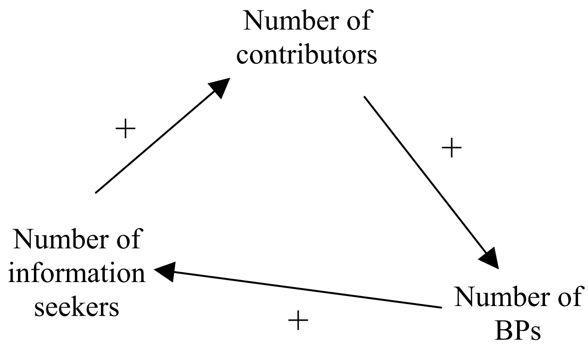

Figure 3. Cause map showing indirect network effects in the SHARE system.

Now, let us turn to the contributors. There may be several reasons why managers in Beta decide to make submissions to the common knowledge repository[3] (Davenport & Prusak 1998). First of all, they may earn status among their peers (repute) (Brown & Duguid 2000). Those who submit good BPs may expect to become known and respected as competent and innovative managers (and such recognition may also lead to career advancement). Second, they may choose to contribute to the common stock of knowledge because they expect that they will benefit from it themselves sometime in the future (reciprocity). Third, they may do it because they want to help others in the company (altruism). Fourth, they may do it simply out of loyalty to their company. No matter what the specific motivations are, it is reasonable to assume that the benefits increase with the number of information seekers and that a large number of information seekers will attract more contributors. This is analogous to what effect the size of the audience has for newspapers, journal publications, and TV broadcasts – as their audience grows they increase in significance. In other words, we have identified a positive feedback loop in the system: When the number of information seekers goes up, it becomes more attractive to make submission and the number of contributors grows. This in turn leads to an increase in the number of BPs, which attracts still more information seekers. See figure 3.

Up until now, we have ignored the question of *quality;* that is we have assumed that the value of the common repository is simply a function of the number of BPs in the database. This is, of course, an over-simplification. In practice, the submissions' quality – in terms of innovation, general usefulness, clarity of presentation, etc. – is of paramount importance. Information seekers are not interested in large amounts of information per se, but in ideas and processes that they can use to improve their own practices. Corporate management at Beta were well aware of this, which is why they established the rigorous review process in the beginning (although it turned out to be a double-edged sword).

Establishing a review process is a balancing act. If, on the one hand, the reviewers are too lax in their assessments and let too many submissions "pass," this is likely to affect the overall quality of the database detrimentally. On the other hand, if the reviewers set the standards too high and reject too

[3] Here, we have assumed that (a sufficiently large number of) people are motivated to contribute to the common repository – if they believe that it will be widely used. However, this may not always be the case. There may be situations where people are unwilling to contribute to the common repository (although everybody agree that the repository would be highly valuable) because it requires too much effort and brings no immediate benefit. Economists refer to such situations as ones with "public goods." In contrast to private goods, public goods are indivisible and non-excludable (Cornes & Sandler 1986): Once the good has been created and is available for one user, it will be available for all users (regardless of whether they have paid for the good or contributed to the production of it.) This gives each individual user the incentive not to contribute in the hope that the other users will provide the good. Relying on the contribution of others in this fashion, the individual is "free riding." If all free ride, then the collective result is poorer (suboptimal) than the individual users prefer. Rafaeli and LaRose (1993) have analyzed the use of computer bulleting boards from a "public goods" perspective.

many submissions, they may well scare off potential future contributors. Thus, the review process may have a negative "side-effect" on the flow of new submissions and may, in the worst case, turn a "virtuous cycle" of growth in to a "vicious cycle" of collapse. See figure 4.

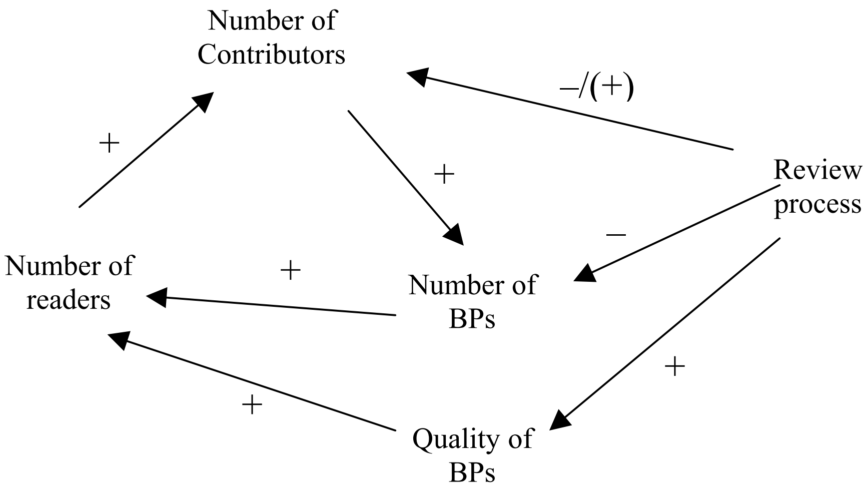

Figure 4. Cause map of knowledge sharing with review process

This is not to imply that introducing a review process will always have a negative impact on the number of contributors. On the contrary, we believe that if the reviewers "get it right" (that is, if they find the right level of acceptance/rejection), then the review process may, in fact, attract more contributors because the prestige associated with "publishing" in the database will be higher (just like prestigious scientific journals attract many authors). In this case, the review process will not only increase the quality of the BPs in the database, but also the number of new submissions.

5.2 Network effects at work

Let us now turn back to the course of events, from the early introduction of the database to the admission of failure and closedown three years later, and take a closer look at the dynamics of adoption in light of the identified network effects.

Phase I. The introduction of the database went wrong right from the beginning, because of the design of the review process. Beta's management was – quite justifiably – very focused on establishing a quality database, but the way in which the centralized review process was effectuated turned out to be a serious mistake. By rejecting more than half of the submitted BP proposals during the first round of reviews the review committee immediately put an end to the flow of new submissions. And without a steady flow of new BPs into the database, it was worthless.

Phase II. The reaction by corporate management was to give up the review process altogether and make it free for all to put new BPs into the database. This change of policy did not, however, immediately change the situation. The managers in Beta had not forgotten what had happened during phase I and they were still not eager to "stick their neck out" and submit their best practices to the database. The SHARE project manager had to "get on his hands and knees" to persuade the managers to contribute to building the database. After a slow start, the project manager's efforts appeared to work: From April until December 1997 the number of new submissions increased as more and more managers began to put new BPs into the repository.

However in January 1998, this positive trend was suddenly reversed, and the number of new submissions started to fall. By December 1998 there were only submitted 9 new BPs to the repository as opposed to 36 the previous year. This negative turn can be explained by the fact that during this period it became increasingly clear to the authors that the number of managers seeking help and advice from the common repository was close to zero. In other words, it became more and more obvious that the database had no "audience." Even worse, the expectation that this audience would materialize some time in the future and that the database would then become an important medium for the exchange of knowledge among managers in Beta also began to fade. Therefore, people lost interest and stopped contributing to the content of the database.

Why were the organizational members not interested in using the database to seek advice and inspiration to improve their own practices? Apparently, the main reason was a perceived lack of high-quality content. In many cases, managers explained that they had browsed the database a few times without finding anything of interest – and then they had forgotten all about it and never went back. Other managers explained that they had heard from others that the database was not worth their time and effort – so they never visited the database themselves. Thus, it seems as if Beta's managers very quickly developed a shared negative opinion about the value and usefulness of the SHARE system – often without having visited it. We see here how general expectations and the opinion of other managers play a significant role in creating a situation where no one wants to spend their time on the system. By the end of phase II, the system is trapped in a *vicious cycle* where the database fails to attract information seekers because of a lack of valuable content, and where the lack of readers and the negative expectations in general, make the contributors loose confidence in the system and give it up.

Phase III. Corporate management did, however, not intend to give up the project. In an effort to save the database, it was decided to redesign the user interface and to establish a new review process to improve the quality of the content of the database. This time great care was taken to design the review process in a way that did not intimidate potential contributors. The new review process was decentralized and based on the principle of peer review.

None of these changes had any effect. Despite a vigorous campaign to promote the "new" SHARE system in the spring of 1999, users stayed away. Very few new BPs were submitted for review and the number of people seeking information in the database stayed very, very low. In our view, the managers in Beta had at this point in time already made up their mind about the SHARE system. They considered it to be a failure, an innovation that would not "take off", and they had little or no faith in the measures taken to change the situation. As a consequence they decided to "wait and see" what others would do, leading to a situation where everyone is "watching the group" – the other managers in Beta – to discern what the group choice may be. In effect, allowing for nothing to happen. In this way, the expectation that the "new" SHARE system would be a failure (just like the "old" system) became a self-fulfilling prophecy. The system never reached critical mass and had to be closed down.

6 IMPLICATIONS FOR IS THEORY AND PRACTICE

Our analysis has important implications for both research and practice. By calling attention to the existence and significance of network effects, this study complements existing research on IS implementation and provides practitioners with useful insights on implementing technologies to support collaboration and knowledge sharing in organizations.

5.2 The start-up problem

The presence of significant network effects makes it difficult to implement new KMS in organizations because the network effects create a "start-up"-problem – similar to the problem encountered by companies introducing a new product or service to the market. The start-up problem refers to the costs and practical difficulties of attaining a viable user community starting from zero (Rohlfs 1974). The trouble is that early adopters experience few benefits (and sometimes high costs) from using the new technology because the number of users is too low. Thus, the adoption and use of the technology, even though useful and beneficial, cannot get started by itself (Rohlfs 1974). It requires some positive action by the implementor – for instance the provision of rewards to the early users. The challenge is to get past the "critical mass" point after which the rate of adoption becomes self-sustaining.

5.3 The role of expectations and self-fulfilling prophecies

Expectations are a key factor in organizational members' decisions about whether or not to adopt a new KMS. People will not invest the time and effort in learning how to use a new system unless they believe that usage will become widespread (Lou et al., 2000). The implication is that self-fulfilling prophecies are inevitable. If people think that the new system will become a success, a bandwagon will form, a virtuous cycle will begin and their expectations will prove correct. But if people expect the system to flop, the adoption process will lack momentum, the system will fail to reach critical mass, and again their expectations will prove correct. As Shapiro and Varian (1999, p. 181) say, "the implication is frightening:" success and failure are driven as much by user expectations as by the underlying value of the system. The SHARE database is a good example of this.

5.4 Measures to overcome the start-up problem

We conclude by briefly discussing a number of practical issues and implementation tactics having to do with overcoming the start-up problem and attaining a critical mass of users when introducing a new KMS in an organization.

First of all, it is important to understand that half-hearted initiatives and half measures are "worse than useless" (Rohlfs 1974). If critical mass is not achieved, the whole effort will be a complete failure and the resources invested in developing and implementing the system will be lost.

Second, it is essential that the system is deliberately designed and implemented to minimize the costs (e.g. as measured in their time) and maximize the potential benefits for early adopters. For instance, it is important that the user interface is well designed and that an efficient support organization and help desk is in place right from the start. (This is, of course, always important – but in the presence of strong network effects, it becomes crucial.)

Third, there is the issue of incentives. In some cases, it may be tempting to simply mandate system use and punish those who do not comply – but in general, we think that this is a poor solution to situations comparable to SHARE's introduction. The use of negative sanctions may provoke resistance toward the system and there is no guarantee that forced participation will lead to the desired results in terms of organizational performance or improved productivity. When people are forced to use a system, they frequently use it in ways that do not benefit the organization (Markus & Keil 1994). Another, more promising strategy, is to provide positive incentives or rewards to the early users of the system. It could be some kind of economic benefits or more intangible, symbolic rewards like official recognition (which in turn may lead to career advancements). However, positive incentive schemes have a backside as well, because of overemphasis on the rewarded behavior. If Beta, for instance, had offered to pay for submissions to the SHARE database, it might have induced people to focus on quantity rather than quality in making submissions (Brown & Duguid 2000).

Finally, one has to consider how to identify the initial target group of users. The question is whether to target the entire community of potential users right from the start or to try and identify a smaller subset of users to begin with. Instead of attempting corporate-wide implementation, Beta might, for instance, have limited the initial use of the SHARE database to a single business function (e.g. production, marketing, R&D) or to a single division and then later – after having gained more experience with the technology – expanded usage to the whole corporation. Beginning with a smaller – and usually also more homogeneous – subset of users may greatly reduce the practical difficulty of starting up the system. There are three reasons for this: First, when the user community is smaller and more homogeneous, it is easier for the system designers to attend to the special needs and interests of individual users and obtain a good "fit" between the system and the organizational context. Second, it allows the implementors to focus on one specific part of the organization and thus take local circumstances and values into consideration when planning how to "market" the system so as to entice people into using it. Third, if the first implementation of the system is a success, then this may well have a positive influence on user expectations in other parts of the organization and, thus, pave the way for the system's subsequent implementations. There is one additional advantage to this strategy. If implementation fails – despite careful planning and concerted efforts to reach a critical mass of users – the defeat will be less devastating and the losses, both in terms of invested resources and prestige, more limited.

7 CONCLUSION

For decades, economists have studied the implications of network effects for market dynamics and competitive strategy in information and communication technologies (Shapiro & Varian 1999). Network effects, however, do not operate in the market place only. They also operate within the boundaries of organizations where they can have a strong impact on adoption dynamics when new information and communication systems are introduced.

The role of network effects in the implementation of information and communication systems in organizations has largely been overlooked by IS researchers. However, many information and communication technologies (e.g. KMS) exhibit significant network effects, which act as a potent force during implementation by creating a "start-up problem" when the new technology is introduced in the organization and by complicating the implementation process in multiple ways (e.g. by making self-fulfilling prophecies inevitable).

We believe that the concepts of (direct and indirect) network effects and positive feedback offer a useful starting point for examining implementation processes, adoption dynamics and use patterns associated with advanced computer-mediated communication technologies in organizations. We also believe that the concept of network effects can be used to diagnose, explain and (hopefully) prevent implementation failures, and hence be particularly useful to system developers and implementors attempting to manage the introduction of technologies such as knowledge management systems.

ACKNOWLEDGEMENTS

We would like to thank the members of the Beta Corporation for their willingness to participate in this research. We would also like to acknowledge the support of the Danish Research Councils (grant no. 56-00-0346 and grant no. 99-00-092). J.P. Bansler gratefully acknowledges the hospitality of the CIS department at Georgia State University where he was a visiting professor in 2001 when the first draft of this article was written up.

REFERENCES

Alavi, M. & D.E. Leidner (1999). Knowledge management systems: issues, challenges and benefits. *Communications of the Association for Information Systems* 1, article 7.

Bansler, J. & E. Havn (2001). Sharing best practices: an empirical study of IT-support for knowledge sharing. *Proceedings of the 9th European Conference on Information Systems,* Bled, Slovenia, June 27-29, 2001.

Brown, J.S. & P. Duguid (2000). Balancing act: how to capture knowledge without killing it. *Harvard Business Review* May-June 2000: 73-80.

Cornes, R. & T. Sandler (1986). *The Theory of Externalities, Public Goods and Club Goods.* Cambridge University Press.

Davenport, T.H. & L. Prusak (1998). *Working Knowledge.* Harvard Business School Press.

Economides, N. (1996). The Economics of Networks. *International Journal of Industrial Organization* 14 (2).

Economides, N. & C. Himmelberg (1995). Critical Mass and Network Evolution in Telecommunications. In: G. Brock (ed.). *Toward a Competitive Telecommunications Industry.*

Katz, M.L. & C. Shapiro (1985). Network externalities, competition, and compatibility. *The American Economic Review* 75 (3): 424-440.

Katz, M.L. & C. Shapiro (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy* 94 (4).

Lai, V.S. & R.K. Mahapatra (1997). Exploring the research in information technology implementation. *Information & Management* 32: 187-201.

Lou, H., W. Luo & D. Strong (2000). Perceived critical mass effect on groupware acceptance. *European Journal of Information Systems* 9: 91-103.

Markus, M.L. (1987). Toward a "critical mass" theory of interactive media: universal access, interdependence and diffusion. *Communication Research* 14: 491-511.

Markus, M.L. (1990). Toward a "critical mass" theory of interactive media. In: J. Fulk & C. Steinfeld (eds.). *Organizations and Communication Technology.* Sage Publications.

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. & M. Keil (1994). If we build it, they will come: designing information systems that people want to use. *Sloan Management Review* Summer 1994: 11-25.

Orlikowski, W.J. (1992). Learning from Notes: organizational issues in groupware implementation. *Proceedings of the Conference on Computer-Supported Cooperative Work,* Toronto, Canada, 1992.

Rafaeli, S. & R. LaRose (1993). Electronic bulletin boards and "public goods" explanations of collaborative mass media. *Communications Research* 20 (2): 277-297.

Rice, R.E. (1990). Computer-mediated communication system network data: theoretical concerns and empirical examples. *International Journal of Man-Machine Studies* 30: 1-21.

Rogers, E.M. (1990). The "critical mass" in the diffusion of interactive technologies. In: M. Carnevale, M. Lucertini & S. Nicosia (eds.). *Modeling the Innovation: Communications, Automation and Information Systems.* Elsevier Science Publishers B.V.

Rohlfs, J. (1974). A theory of interdependent demand for a communication service. *The Bell Journal of Economics and Management Science* 5 (1): 16-37.

Shapiro, C. & H.R. Varian (1999). *Information Rules.* Harvard Business School Press.

Stake, R. E. (2000). Case Studies. In *Handbook of Qualitative Research.* 2 ed. Edited by Norman K. Denzin and Yvonna S. Lincoln. Sage Publications, London.