PART 4

Knowledge Management and Innovation

In this section, the complexities involved with knowledge management, information sharing, and innovation will be discussed. Advances in computing and information technology have made possible the accumulation of vast amounts of data, leaving users with the problem of organizing it efficiently. The field of knowledge management has attempted to solve this issue through the discovery and organization of relationships within data, and the identification and documentation of rules for managing it. The following chapters provide insight into improved methods in the valuation of R&D, the importance of the proper utilization and management of information, and the steps corporate management can take to foster innovation.

In Chapter 11, Larry Chorn of the American Graduate School of International Management provides a valuation method for determining the financial value of investments in knowledge. While previous attempts based on discounted cash flow have encountered failures, Chorn believes that his "Real Options analysis" overcomes the difficulties, and enables a correct valuation of returns on knowledge investment. He provides examples of successful valuation of a pharmaceutical R&D program and multinational energy companies to illustrate how Real Options analysis can aid management and shareholders in their decisions to invest in R&D. As expertise with his technique increases, Chorn believes "more and more investments in knowledge and knowledge management will be justified financially."

Stephen Younger of the Los Alamos National Laboratory (LANL) details the importance of information in nuclear deterrence in Chapter 12. He argues that information is a strategic asset, and explains the role that it plays in maintaining the nuclear arsenal. Since the signing of the Comprehensive Test Ban Treaty, the United States is no longer actively designing or producing nuclear weapons. This makes the existing information important for both the utilization of current knowledge and the generation of new knowledge for the program to "fill in the gaps in our fundamental understanding of the physics of the weapons." Younger believes that the generation of knowledge provides an economic asset, and notes that the research done by LANL, and the other national laboratories, has been at the forefront of diagnostic and computing technologies, and has generated economic value for the U.S.

Stephen Hodges of IBM discusses in Chapter 13 the manner in which knowledge is accumulated, managed, and ultimately shared in industry. Hodges asserts that knowledge management is "the capability to search, capture, retrieve, share information, evaluate, and transfer." As we become better at accumulating knowledge, and simulations and models become more complicated, better human-to-computer interfaces are needed. He predicts that virtual reality will fill this role, with people connected in "not just a one- or two-dimensional interface via terminal to computer, but a three- and four-dimensional interface." This expanse of data will require knowledge management solutions. However, in order to define a knowledge management system, Hodges insists that groups must first answer a number of important questions: "What is the work product? What is its purpose? What is the impact of not having it?"

Michael Laird continues the discussion of knowledge management and knowledge sharing in industry with a perspective from the Xerox Corporation. In Chapter 14, he explains the concepts behind knowledge sharing and information management, and the difficulty in choosing a successful strategy. Information management is concerned with "capture, retention, and standardization," while knowledge management, "a major enabler for organization effectiveness," is a movement within organizations to organize assets, technology, and people. He believes the key distinction between the two is that knowledge management relies on human input, "which is the opposite from standardization." Laird asserts that creating a proper knowledge management project is difficult, and requires a project strategy if success is to be assured.

In Chapter 15, Richard Leifer and Mark Rice, both of the Lally School of Management and Technology at Rensselaer Polytechnic Institute, discuss the difficulties managers must face when determining whether to attempt incremental or breakthrough innovation. The authors offer to "provide managers with a starting point for experimentation with new practices and tactics for improving breakthrough innovation effectiveness and to provide researchers with fruitful directions for further study." Leifer and Rice based their research on a four-year study of twelve breakthrough innovation projects. Their results show that breakthrough innovation can be fostered through five implementation tactics: stimulating attractive ideas; promoting opportunity recognition; evaluating and screening discontinuous innovations; creating incubating organizational structures; and catalyzing individual initiative. They conclude that their tactics can be used to balance mainstream business activity and breakthrough innovation activity to enable future development.

11 Managing Corporate Knowledge to Create Strategic Opportunities

Larry G. Chorn

Evidence exists showing that markets and investors reward corporations that manage and act on corporate expertise to create shareholder value. I intend to show you that it is possible to measure the market's valuation of corporate expertise. The implication from this observation is that markets expect corporations to use that in-house capability to add value to assets and grow their business revenue. The example used to illustrate this observation comes from the energy industry.

Second, I will use a simple example from the pharmaceutical industry to illustrate the application of corporate expertise to product creation. Many of us have been searching for a financial tool that will quantify the intangible value of technical know-how and research activities. My example will apply a new valuation and decision management technique that some of you may have already seen in the finance and business literature, Real Options analysis. Not only can we use it to value the creation and application of knowledge, but also to value knowledge management activities and the efforts of knowledge managers.

Introduction

We are all aware of the difficulties and failures of traditional valuation techniques based on discounted cash flow (DCF). These techniques are unsuitable for valuing risky opportunities in which success is based on the arrival and application of new knowledge. The techniques re-

---

Larry G. Chorn is a visiting professor of finance at the American Graduate School of International Management. *This article is based on remarks delivered at the 24th Annual AAAS Colloquium on Science and Technology Policy, held April 14–16, 1999, in Washington, DC.*

quire the use of chance factors for success that are nearly zero, reflecting the risky nature of the investment. Other investments that require long time frames between the initiation of a knowledge-driven program to its successful implementation are penalized heavily based on the time value of money concepts in DCF. As a consequence we have become accustomed to hearing that investments in knowledge have little or no financial value.

But we know that corporations routinely ignore the guidance of traditional value techniques. There are many examples of investments in opportunities that provide flexibility for managers to respond to the arrival of new knowledge about the project, the implementation of new technologies in the project, and new market needs. Clearly these investments reflect corporations' instinctive valuation of growth opportunities and recognition of the embedded ability to manage the project's risk.

To exploit these opportunities and manage the associated risks, corporations must sustain and manage business and technical expertise so that it is readily available for application. Growth opportunities depend on creating new knowledge, exploiting existing knowledge in new methods, and protecting hard won knowledge from exploitation by competitors.

These efforts must be measured in terms of new products, ability to add value to existing assets, and growth of revenue and profits. Here is where the traditional valuation techniques fail to support our instinctual insights. Applying financial tools that give a more realistic representation of the business processes of risk management through the application of knowledge management will show the value we believe has always been at hand but was never quantifiable.

Market Rewards for Corporate Knowledge Application

The first example we will consider looks at how the availability of corporate technical expertise in developing new pharmaceuticals can be incorporated into a Real Options analysis of a research and development (R&D) project. In this example, we will see how improvement in valuation can make the justification of a good project much easier.

There is evidence that markets reward companies for their ability to sustain and apply technical expertise to growth opportunities. In turn, these opportunities lead to the traditional measures of shareholder value. I will demonstrate one example of this reward process using an example of a multinational energy company's management of its oil and gas production portfolio, and its share price on the New York Stock Exchange.

Applying Knowledge Management to Identify High Value R&D Opportunities

Pharmaceutical R&D is driven by the company's exploitation of its scientific expertise. New compounds are tested for their curative potential against a disease over several clinical stages. Each stage has a high risk of failure. Taking the pharmaceutical from the discovery stage to commercialization requires multiple years of costly effort. Later stages build upon prior work and stages cannot be skipped. Ultimately, if a drug survives each stage, the company can earn many years of patent-protected profits in return for its efforts.

As with most R&D investments, the cumulative risk of failure across the many stages to commercialization lowers the likelihood of receiving the profits of a successful commercialization. Even more difficult on the valuation is the extended period between initial, costly R&D investments and the beginning of revenue from sales. Traditional discounted cash flow usually cannot justify the investment in a promising new pharmaceutical. The project's net present value (NPV) will be negative.

It is becoming more and more commonly recognized that stage-wise R&D projects, like the one described above for the development of a new pharmaceutical, can be considered as a collection of option purchases bundled with the acquisition of knowledge. We can view the first stage of the project as the purchase of knowledge about the relationship of drug efficacy versus toxicity or side effects in a small subject group plus the option to invest in the next stage of development, which is larger scale tests. The option gives the company the right, but not the obligation, to undertake the next stage. If the first stage is a failure, the company is not required to purchase (invest in) the next stages. However, the company cannot undertake the second stage without purchasing the knowledge that comes with completion of stage one.

Consider a simplified example built from a valuation study that I recently conducted for a major U.S. pharmaceutical manufacturer. The company believes that a new compound has the potential to treat a significant neurological illness in older patients. If this compound can pass through the four stages of testing and approval, its worldwide market potential is approximately $500 million in after tax earnings over its patent life, discounted back to the year of its initial introduction. The discounting of revenues each year in the future uses the company's weight-averaged-cost-of-capital as a discount rate.

There is considerable uncertainty in the pharmacologists' minds regarding the potential for success in the early stages of study. As the compound successfully moves from one stage to the next, the overall risk for the project decreases. However, the cumulative risk is so high that a traditional financial analysis of the project indicates a negative return on investment. The NPV is negative. In spite of this unattractive financial metric, pharmaceutical R&D managers routinely undertake investments like this. They disregard the existing financial indicators and invest based on their instinct and experience in the industry.

Management makes a decision to invest or overlook the pharmaceutical's potential in phase one. The outcome of phase one is either good or poor. This leads to the next decision, to invest in phase two if the outcome of phase one was good, or to abandon the project if the outcome was poor. This invest-learn-evaluate-decide cycle continues for each of the four phases in the pharmaceutical's development. The manager has the right, but not the obligation, to invest in the next stage, only if the prior stage was purchased.

How does corporate knowledge and its management enter into the valuation of the option and guide management's decision to invest in the R&D project? An option is the purchase of a period of time to learn more about the value (potential) of the underlying asset, which is, in this case, the new pharmaceutical. The uncertainty around the value must be quantified and then resolved during the time that the option is "alive." If the resolution is completed and the result is favorable, management is guided to exercise the option, i.e., purchase the next phase of the R&D program. If it is unfavorable, the option expires unexercised and the program is abandoned. Similarly, if the uncertainty is not resolved within the option's life, the guidance is to abandon the program. Project experts must provide a quantification of the uncertainty of each investment phase. Their insight into the initial degree of uncertainty and how that uncertainty will be resolved over time is the critical input to the option valuation. By drawing upon project experts and corporate knowhow archived in data warehouses, we can value the R&D phase relative to the amount of risk it removes from the project's outcome.

While the results are not guaranteed, we have a measure of the value of attempting to remove uncertainty, and that is what Real Option analysis quantifies. Consequently, we find that technical expertise and the ability to create and manage knowledge adds value to an investment opportunity. This is just the financial tool that we have been looking for!

But how much value does a Real Option analysis reveal for this three phase pharmaceutical R&D project? If the DCF analysis was negative for both phases one and two, then the value of the project does not become positive until stages one and two have removed enough risk that the NPV becomes positive. If management follows the standard investment rule, neither of these phases will be funded and phase three would have never occurred.

The Real Options analysis, on the other hand, indicates that all phases have a positive value. The value estimated for phases one and two is predominantly option value. The knowledge gathered in these phases was not of high enough value to independently result in a positive return on investment. But the option for further investment, if phases one and two were successful, was quite valuable. By phase three most of the uncertainty in the project has been resolved by knowledge creation in phases one and two. Consequently, most of the option value has been converted to tangible value. Using Real Options analysis, management would have been guided to invest in phase one and two with the expectation of receiving value in excess of its cost.

Adding Value Through Technical Know-how

Multinational energy companies maintain large staff groups skilled in the technologies of exploration, development, and production of oil and gas. These groups represent several hundred years of combined experience in maximizing exploration efficiency, optimizing development designs for lowest cost, increasing safety in production operations, and increasing recovery efficiencies from oil and gas reservoirs. The companies' economic success is tied to these staffs' capabilities.

Each year an energy company produces a significant fraction of their hydrocarbon reserves inventory. If these assets are not replenished the company is effectively in liquidation. It will go out of business within the decade. Consequently, shareholders pay very close attention to discoveries and reserve replacement reports from each company.

New reserves create flexibility and the opportunity to exploit the company's technical expertise in developing and producing hydrocarbons as efficiently as possible. If the company's business skills are inadequate to capture the opportunities to explore and develop, the expertise is wasted and atrophies away through lack of use, resignations, and retirements. Adding new reserves creates the opportunity to:

- discover additional hydrocarbon reserves in adjacent geologic structures,

- increase production from those reserves through the application of new drilling and production techniques, and

- lower production costs by sharing expensive production facilities among several fields.

The addition of new reserves leads to more opportunities and amplifies value creation.

If we look at information provided in annual reports by the multinationals, we can correlate their reserve additions, value creation, and share price over a multi-year period. I have done this for twelve of the largest companies and found a repeating relationship between the three measures. Further, there is an important correlation between reserve addition and share price.

If we were to create a plot of reserve replacement (y-axis) and inventory value change (x-axis), we would see the axes cross at 100 percent reserve replacement and zero percent change in reserve inventory value. Annual performance by the company that replaces more than 100 percent of the reserves it produced and sold is favorable. Below 100 percent represents liquidation. Similarly, a positive change in inventory value is favorable to shareholder wealth creation, and a negative change is unfavorable. This suggests that companies will strive to be in the upper right hand quadrant—growing inventory reserves and value. Outcomes in the lower left quadrant are unfavorable representing a poor performance by management.

Note that in the expansion phase it was creating future opportunities but reducing the value of its portfolio by adding lower valued assets. In the exploitation phase it was adding value to the inventory in spite of its shrinkage. Application of technical knowledge was increasing producible reserve volumes and lowering costs. Reiterating, in the late 1980s, the company was spending capital to create the options for application of its technical expertise. In the early 1990s, it was exploiting those options and adding value to the portfolio with higher valued assets. By the mid-1990s, the company was again spending capital to create new options by expanding its inventory.

We would expect the market to reward the value creation stage, i.e., the early 1990s. But data suggests the opposite is true. The market rewarded the company for creating opportunity to add value through its technical expertise. During the early 1990s, when value was being added to the portfolio, the market was neutral to those events.

My interpretation of this is that the market expects companies to sustain their production position within the competitive framework of the industry. Reserve replacement and a gradual growth of value in the asset base is expected, not rewarded. When a company's management creates opportunity for its organization to apply technical and business expertise to increase the value of existing assets, the market responds with a higher valuation. This company has highly regarded technology and excellent business knowledge; it needs the opportunities to apply those skills. When management acquires the opportunities, the embedded knowledge of the corporation can create value for the shareholder, and the market responds accordingly.

Adding lower valued reserves did not diminish the company's value in the eyes of the market, but it did create opportunities for the company to grow. This reflects the mission of a corporation with an extensive technical and business knowledge base: to take lower valued assets, apply technical resources and add value to the asset, thus creating additional new opportunities.

Conclusion

While we all intuitively recognize the value of knowledge creation and management, there has been little validation of that intuition through financial metrics. Since so much of our R&D and knowledge management activities reside in the corporate domain, this lack of financial quantification has strained the willingness of management and shareholders to invest in these activities.

We have demonstrated that viewing investments in knowledge and knowledge management as the purchase of options on future revenues leads to increased valuation and new revenues for a corporation through two examples in this presentation. Quantifying the value of this option creation can be done through the application of Real Options analysis. As expertise with this technique increases, I believe more and more investments in knowledge and knowledge management will be justified financially. Corporations and governments will use this approach to identify valuable opportunities to create value for their constituents.

12 Knowledge Management at Los Alamos

Stephen M. Younger

This article will focus on the role of information as a strategic asset and as a component of our nuclear deterrent (in ways that most people have not thought about).

The core mission of Los Alamos is to maintain America's nuclear deterrent, a mission that has remained remarkably constant through the end of the Cold War. Nuclear weapons are in the words of every president from Truman to Clinton, in the supreme national interest of the United States.

Everything associated with a nuclear weapon is described by a superlative. They are the "hottest," the "dense-ist," and the "most powerful" weapons. They are instruments that make other countries deal with the United States in a very cautious manner. People tend not to think about other countries invading Florida. Although many times in this century countries have invaded one another, America's nuclear arsenal dissuades or deters countries from invading us. That deterrence is expected to continue. Officials of the United States government have said that they expect nuclear weapons to play a role in national security for some time to come, even though we are actively pursuing arms control treaties (when they are in our national interest, of course).

That core mission has remained constant, but the way in which we have pursued that mission has changed in at least two very fundamental ways. We originally developed nuclear weapons the same way people developed everything from toasters to rocket ships—through a process of design, testing, and production. But now, because the President has not identified a need for a new nuclear weapon design at this

---

Stephen M. Younger is the associate director, Office of the Associate Laboratory Director for Nuclear Weapons, Los Alamos National Laboratory, New Mexico. *This article is based on remarks delivered at the 24th Annual AAAS Colloquium on Science and Technology Policy, held April 14–16, 1999, in Washington, DC.*

time, we are not actively pursuing the design of a new system. Also, because the President has signed the Comprehensive Test Ban Treaty, we are not actively producing nuclear weapons. As a matter of fact, the production line is in full reverse. We are taking weapons apart, as are Russia and some other countries.

So we have moved from a design, test, produce to an information-based or a knowledge-based method of maintaining the nuclear arsenal.

Why do we need to do anything at all? Nothing lasts forever. You would not put your car in the garage for 20 years and then say, "I'm going to bet the Nation on starting on the first turnover." Cars age. Materials need to be replaced and batteries wear out. You need to do all sorts of things. Each component has a finite life, although with proper attention the whole system may last for a long time. (We have the same situation with nuclear weapons).

I see four different aspects of knowledge. The first aspect is the utilization of existing knowledge. This is especially important when you can not get any more. The second aspect is the generation of new knowledge in certain aspects of the program. The third aspect is information as an instrument of deterrence. The fourth one is information as a generator of economic value, in addition to the value it has for national security.

In terms of utilization of existing information we are in an interesting position. When the Department of Energy downsized its production complex, Los Alamos National Laboratory inherited 25 million records about the weapons that had been built and are currently on alert. The challenge is that when you have 25 million of something it is almost as bad as having none of something. The government measures records in terms of cubic feet (I forget how many cubic feet we have). It almost gets to the point where you have a bulldozer pushing these records into a little corner and then someone asks, "Can I have the inspection data for this serial number part that was made in 1978?" Try to find it. Does it even exist? Can you read it? Is it on media you can deal with? Those are the challenges.

We also have data from 1,000 underground nuclear tests and some atmospheric tests. These data are irreplaceable since the United States is no longer conducting nuclear tests. We are not generating new information on the actual performance of weapons. This raises secondary questions of how we train new people who will never have the experience of doing nuclear tests.

The second aspect of knowledge, the generation of new knowledge relevant to nuclear weapons, includes filling in the gaps in our fundamental understanding of the physics of weapons. For a long time we used this design, test, produce mode. When you use that sequence you do not necessarily have to know everything there is to know about a subject. For example, if you are building automobiles you have to know, in general, how the metal behaves and so forth but you do not necessarily have to have the absolute best understanding at the atomic scale of how the fender is going to perform over time.

On the other hand if you are not permitted to do tests then the importance of that fundamental understanding increases tremendously. The whole basis behind the Department of Energy's Science-Based Stockpile Stewardship Program is that your confidence should go up as your knowledge goes up. The more you know the more confident you should be. So the generation of new and important information, new detailed data about things like plutonium and high explosives, is critically important to the future of the program. We are following the same kind of process that any scientific enterprise would use—a combination of theory, computation, and experiment. Computation now occupies that middle ground between theory and experiment. The role of large-scale computation has become considerably more important for us. The reason is very simple: If you cannot do a nuclear test, the only way to estimate the safety and reliability of a nuclear weapon is through very large-scale computations.

Los Alamos currently has the world's fastest computer at 1.608 tera-ops with a peak capability of three tera-ops. That is three trillion operations per second, with multi-terabytes of storage. These are gigantic machines. But they are not yet able to do first-principles calculations of what we need. They will be replaced in time by a ten tera-op machine at Livermore, a 30 at Los Alamos, and so forth. These machines open up tremendous scientific vistas. Some significant scientific challenges are associated with just using the data that come out of these machines.

Files that are multi-terabyte in size exceed the capability of a human being to absorb the data or information. You will not live long enough for your eye to absorb all of that information. So we need to look at how human beings interact with machines. It is a very interesting subject because until now almost the entire interaction has been on the machine's terms. That is, you learn how to use the machine and you type in commands or you click buttons or things and the machine tells you how to interact with it.

In the future we are not going to be able to do that. We are going to interact with the machine on human terms, using as many of our senses as possible. The machine is going to have to understand what we need and what we want. We may need to go as far afield as looking at art versus mere data management. For example a painting may not be as accurate as a photograph but it conveys the impression. High resolution is not the only thing required to learn what you need to from these very, very complex calculations.

The third aspect of knowledge is information as an instrument of deterrence and its role in defense. When the United States was conducting underground nuclear tests, perhaps one of the choicest assignments of foreign intelligence agencies was to go to Las Vegas, put a seismometer down, wait for the test, and watch the ground shake. It was clear to everybody that the United States was capable of producing significant yields in nuclear explosions.

We are not doing that now. So what is to prevent an adversary 15 years from now, after reading a series of articles in the *Washington Post* exposing how we are not taking care of America's strategic weapons and everything is turning sour, from taking his chance? How do we prevent him from doing that? We prevent him from doing that by projecting confidence in our abilities in as many areas as we can. We do that by publications and conferences. Having people from the weapons laboratories out in the community constitutes an element of risk from a security standpoint, but it is absolutely essential from a deterrence standpoint. If people from Los Alamos, Livermore, and Sandia are out there giving superb papers in science and engineering, hopefully the scientists in a (less-than-friendly) nation will go back and tell their leader about it. "These people are really good in what they tell us, and I can see the connection to nuclear defense. Perhaps they are also good on the other side of the fence." That is an important part of how we need to do our business in the future.

This idea also leaps over into conventional defense. A recent report questioned the utility of air strikes against Scud missile launchers in the Gulf War. Apparently Saddam Hussein set up a lot of fake launchers. They had crews and radio signatures that were associated with Scud launchers. The real ones never talked to anybody on the radio and were carefully hidden. Here the issue was not one of firepower. It does not take a very big explosion to destroy a missile. The critical issue was knowledge.

The fourth aspect of knowledge is the generation of information as an economic asset. There is enormous value to the Nation in the information resident at the Nuclear Weapons Laboratories. An illustration of that is the fact that a number of countries spend significant amounts of money trying to get that information. You can read about some of that in the newspapers. The knowledge is so valuable that when a problem is perceived it becomes an urgent national issue. The President becomes engaged and congressional committees become engaged.

But the laboratories do a number of other things, too, and these things generate economic value. We have always been on the forefront of a variety of diagnostic technologies and computing technologies. Economic value is generated by the activities of the laboratories and that must be protected as an investment of the taxpayers. We must also make sure that these advances benefit the United States rather than other countries. The reason for that is quite simple: American taxpayers supported it so American taxpayers should derive the value from it.

Knowledge is critical in maintaining the deterrence and keeping other governments confident that America's weapons will work if called upon to do so. It is important in how we structure our programs to add to our knowledge base in an ever resource-limited environment. And it is important in understanding how our activities create economic value for America.

13 Directions for Collaborative Scientific Dialog

Stephen D. Hodges

Knowledge—information—accumulates in stages. A group of scientists, for example, might start with an initiative and some ideas that they share. Over time, they would develop that into new information or new knowledge, which I will label a competency. At that point, if what they have accumulated is of value, we would like them to formally codify it into methods, approaches, and techniques—and begin to teach. So at this stage we call it "competency leadership" because these are the people teaching other scientists or other knowledge workers about what they have discovered and codified. Finally, this would mature into a broad base of networked communities sharing a body of knowledge with each other based upon the course of that accumulation over time.

Today all of the technology companies—including IBM—have some rudimentary capabilities to support the accumulation of knowledge. These might include enhancing infrastructure, such as with an e-mail system. Many people trade information that way—sometimes attaching detailed bodies of thought. There comes a point in that behavior, though, where you have to start thinking about what is and is not confidential, what has a cost, and other considerations. Security issues, for example, apply to basic infrastructure as well as to more complex technologies.

Infrastructure enhancements therefore enable the accumulation of knowledge in that first phase. You can find documents, scan them,

---

Stephen D. Hodges is the managing principal and global competency segment executive for IBM's Work Flow Knowledge Network and New Solution Consulting Group. *This article is based on remarks delivered at the 24th Annual AAAS Colloquium on Science and Technology Policy, held April 14–16, 1999, in Washington, DC.*

index them, and perhaps put them through engines that will glean certain aspects and put those in sophisticated indexes—allowing you to get at them without reading the whole thing. You then can move to work management—essentially, moving packets of information around an organization, or among scientists, in connection with the work that they do. It implies that we have to first understand the work that they do as well as the sequence and alternate sequences in which they do the work. Otherwise, we are not able to move the information with them. We require them to go get it at every point. This is not a new technology.

Knowledge management, however—the capability to search, capture, retrieve, share information, evaluate, and transfer—is fairly new. You encounter this capability on the Internet today, where you can use text-based search engines to find things. The newer search engines, however, are not thinking "text" any more. Instead, they are keeping track of your personal search history and, from that, are making deductive representations in the software about your interests—persisting out there on the network while you are sleeping or not at work, and accumulating more information for you based upon your historical search profile. So that technology does exist, and it is going to come into play more and more in the near future.

Finally, we have the issue of very high-powered computing, and the kinds of simulations and things it enables us to do. We still do not have very good interfaces, so it may take an army of programmers and scientists to construct the complex simulation and we still have consoles, terminals, text, and numbers to interface with the computer. That is not good enough. What we will see coming into play within our lifetimes is a more "virtual reality" sort of thing, where you will have a software representation of yourself resident in the computer. You could be connected physically in some sense to the computer, which is then moving your software representation through a virtual reality. The exciting thing about that is the capability it will give us to interact software representation to software representation—not just a one- or two-dimensional interface via terminal to the computer, but a three- and four-dimensional interface to the computer. It will then be able to respond more quickly, and it will take much less time to construct the simulations.

Establishing a Knowledge Management System

When working with groups to establish a knowledge management system, we first define what that group's work products are. We then establish a dependency diagram among them. For example, if we are working with a group of people who are responsible for building something or constructing an experiment, we would ask them about the sequence in which they do things and then chunk those things into work products. At every point, then, they have an internal checkpoint or an internal delivery to the project in the form of a work product, and we establish a relational dependency diagram among them.

We then describe each of the work products: What is the work product? What is its purpose? What is the impact of not having it? If you are not having it, what is the reason? What notation is used during it? We then provide an example of its output, the development approach to creating it, validation and verification that its is correct, and referenced advice and guidance to people who are new to executing the work product. We also describe other scientists or knowledge workers who have used it, and provide some estimate of the amount of time that it takes to execute the work product.

Each work product has these kinds of characteristics. What is different about characterizing them this way is that it is content-oriented rather than project-oriented. The problem of project management meeting knowledge worker or scientist can be solved to some extent by representing the content part in work products.

Next, we group all of the project's work products into domains. The domains of a given project would differ. If a project has a procedural sequence to be described, we assign a process domain—and those work products are in the process domain. An organization domain indicates that after the project, the roles of people must change. There is an application domain, which is where software and similar systems typically are described. If there is a technology being used, we assign a technology domain that describes the work products related to technology creation and its attributes. Finally, there is a management domain, which describes the project management-oriented types of work products.

Knowledge Management Solutions

We find in IBM that knowledge management solutions typically are built using Lotus frameworks, because we purchased Lotus. My group, in consulting, is not handcuffed to them. We are using Microsoft- and Lotus-based knowledge management solutions, and some of the things we are learning relate to interaction in knowledge spaces. For example, we now are able to have e-meetings within my internal group, wherein we connect to a site on the Internet—either a Microsoft- or a Lotus-based site. We can see each other because we have little cameras that connect. We can also look at a "workspace" in which somebody can post a Microsoft Word or PowerPoint "exhibit." There is also a chat capability and—if you add a conference call to that—real-time interaction.

These collaborative spaces allow us to connect project personnel and clients—particularly during the period between the start-up of the initiative and the onset of large-scale computing. It is a fascinating facet of knowledge management from a behavioral standpoint, because it is where all the behaviors you like and do not like occur. It also represents what technology will bring to bear on knowledge management in the next few years.

14 Knowledge Sharing: A Perspective from Xerox—The Document Company

Michael W. Laird

If Dilbert were in the office and the phrase “knowledge management” came up, he might say, “Gee, I think I hear a buzz word.” We all have our own phrases on what we would like to call this area but I would like to call it “knowledge sharing.” Much research has been done to indicate that knowledge sharing is the way knowledge grows. You do not manage knowledge to make it better; you share it. Research also says that knowledge gets bigger when you share it. The real question that Dilbert raises is what is going on here? What is this thing that we are struggling to understand?

I propose that there is a lot going on in this area that is real and substantive, and that it is important to decision making in large organizations, as well as in small work groups. I will first discuss what knowledge sharing is and then give an example of how it works, as well as how to get further information.

Many data sources indicate that the number of knowledge conferences, as well as publications, has greatly increased over the past few years. Market research in the knowledge management area has also increased over the past four or five years. Forecasts show that consulting revenue in this area is growing and it is going to become a multi-billion dollar business. From the Xerox perspective, the forces driving knowledge sharing are in the same category as the total quality management effort

---

Michael W. Laird is the vice president of Xerox Software Solutions Development at the Xerox Corporation. *This article is based on remarks delivered at the 24th Annual AAAS Colloquium on Science and Technology Policy, held April 14–16, 1999, in Washington, DC.*

initiated as a response to Japanese competition in consumer goods and automobiles in the 1980s. In the 1990s, we saw process reengineering as the major initiative. Movement within organizations was generated by many sources including global competition and the recessions of the early 1990s.

Knowledge management is a movement of the same class. It is aimed at managing knowledge assets and, just as the other movements were, it is driven by a lot of factors, including Internet technology and the democratization of how organizations work. This is not purely a technical phenomenon. There is a huge opportunity here for research in understanding what is happening from a social perspective. This kind of research would move this whole phenomenon forward.

We know there is a big difference between data (organized numbers, bits, bytes, etc.) and information (which is more organized, arranged, and processed). But what do we mean by knowledge? At Xerox our perspective is that knowledge includes the things that one has to have in order to launch into productive action. That is different from data or information by itself. When you launch into productive action, three things must be considered.

First, what is this related to? Where did this come from? How is this connected to something else I know or to a person I know? Context typically comes in the form of a document. There is a paper on a subject, or a book, or a speech. The context is very important to taking action. Second, we have to consider authenticity. Can I trust this? That is a key part of knowledge taking action as opposed to just information (which is simply organized and presented). The third piece of this is experience. Have I done this before? Has my co-worker done this before? These three elements of knowledge are not elements of data or information. This is where social science comes in. We are talking about a phenomenon that is very much technically driven, but it is not just technology. A great deal of social interaction is going on. Somewhere, our society is changing because of the forces that are driving knowledge sharing.

If we say that information is different from knowledge then we would have to say that information management is different from knowledge management. Information management has been around for a long time under a variety of names: information technology, MIS, CIO, etc. We all have some kind of understanding of information management. People are interested in delivery, accessibility, and up time. And there is a very heavy technical focus. This is where the alphabet soup comes in.

(Very few of us can understand this and even some of the people who are in the midst of it do not understand.)

Capture, retention, and standardization are important themes within information management. But knowledge management is a bit different. For example, human inputs are very important, which is the opposite from standardization. We have to trust whoever the person is that is associated with that information, rather than trust the standardization and definition of the data or information presented to us.

Something very interesting is going on in this complement of technology and culture and social science. It is leading to actionable content, something that gets filtered, synthesized, interpreted, analyzed, confidence-created, de-risked, etc. We can then finally come to the decision that lets us act. That is the essence of it.

People look at the whole, but then more information comes in. This information may be on a different subject and from a different endeavor, but it is brought into the conversation and it makes an impact. It may cancel something out or add confidence to some other piece. That, in turn, generates more conversations and more documents.

Documents are a repository, a container, for the interaction that occurred in some group. They document the thought process, who said what, and why we do something. That document then becomes a credibility factor in the next conversation. When we talk about knowledge sharing or knowledge management, at least at Xerox, we ask how we can foster these kinds of interactions.

This moves to an illustration. We in Xerox have a knowledge-sharing system that is used within our field service organization, which has about 20,000 people all over the world. They make a million calls per month on customer sites, either for prevention or maintenance. That is an awful lot of knowledge that is being used, refreshed, and refined. And it is a continuous process. We needed a way to tap that information and leverage it. We developed a process supported by some software products, but more importantly by social engineering. A set of awards and a set of roles and responsibilities foster reporting information up to another group that has dealt with the problem before.

We are a printer company. We know about streaks and paper jams. When it comes to paper jams we know more than anybody else does. This is all codified knowledge. We have very large databases that cover the “red streak problem” and the paper jam problem on a DocuTech-135, for example. This knowledge is very well organized and codified. But we then make new products, and customers come up with new uses for current products. People put our products to uses we never imagined. We find new problems even in old products. We also go into new marketplaces where there are cultural differences. Things are changing all the time.

People are rewarded for figuring out how to fix these problems. We then put the solution into a database. This database is very extensive. We have people who specialize in streaking problems across all products and people who specialize in a particular product and the different kinds of problems that can occur with that product. We have created a very matrixed and complex organization.

Having a tip accepted is similar to getting published in academia. Tips come in from the field and are reviewed. The best tips on a particular problem become the way we do things. We publish the fastest, cheapest, and most replicate-able and appropriate methods for that problem in that product. This is a knowledge-sharing system. Anyone who runs into a problem gets the benefit of a million calls a month and all the experience generated by them. Employees do not have to reinvent the wheel every time a problem occurs. Also, they can contribute. What they find out is leveraged to everyone else.

We have seen a whole culture change within our very large organization. In fact, we now have multiple organizations. Those 20,000 people do not report to one person. They report to dozens of people all across our organization. But they all have the same incentive: to be published. They do not get tenure, but they do get incentives.

Our method of knowledge sharing is like doing research. At the bottom you might have research associates in your university or research lab who deal with very unstructured problems. They have the same need to communicate what they find. This is, in fact, how this particular product was created within Xerox. Some people within our research organization wanted a better way to communicate their results in the problem-solving process. They did not want a highly structured system where someone turns the document over to an administrator who puts it into an indexing system. They wanted something they could put information and documents into. Then they could decide who gets it and who does not get it.

Sponsorship is nice. But there is always a question of how much I should share with that person, and how often. Does the sponsor want to know every week, every month? With this system people can check the Web site whenever they want. It is not the responsibility of the research manager anymore to inform different people who are casually interested. People can come and take a look at it. They are entitled to look at the documents that are shared. (Of course, people cannot see things that are not ready to be shared).

You can look at this as a school. The 20,000 people are 20,000 students and they are solving problem sets. The people at the top are pushing problems down on them. The community product leaders are curriculum leaders. The problem sets that come up are matched with how well the person has matched the curriculum.

This is a generalizable situation. Many activities across many different organizations are similar to field service. They all are amenable to leveraging the accumulated knowledge that is within that whole system of people.

We have found out all kinds of things as a result of working with this. What benefits have we seen? The field service organization has been able to validate what goes on. They are identifying who is good at this and they are rewarding them. These people get more time to act as problem solvers and to leverage their expertise. We have also found that sometimes the problem is not with the machine or with the user, or the intended use they are putting it to. The problem is with our documentation. We need documentation that can be searched and organized. A whole cluster of problems came into our documentation group. The writers realized they needed to write the manuals differently. That has strengthened the way we communicate with our users, and it has reduced our costs. The improved documentation has also improved our customers' effectiveness. It is a win-win situation for everyone.

On an even slower feedback loop (but in the long run even more powerful), a variety of design problems have come up. That goes back to the engineering organizations and who can solve the problem in the next version of the product.

Of course, a lot of this has been going on for a long time in business management. You could say it is just good management. Why is this different and why is this getting a name like "knowledge management?" Something is happening to make all this happen more quickly with more people. The old paradigm for management was that if you wanted to do something quickly you wanted a small number of people involved. Now we have tools and social techniques that help us get things done quickly with tens of thousands of people. That has not happened before.

In every country in Europe we have a separate field organization. In France we have all kinds of data and a system to support the collection of data. This system provides statistics that things are getting better.

Figure 1
Characteristics of Knowledge Management Projects

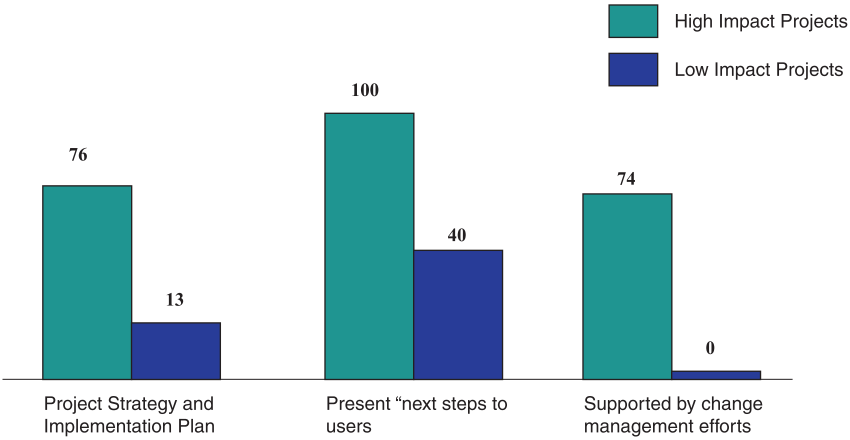

N = 93 projects at 83 companies
Source: Teltech Resource Network Corp.; CIO Enterprise, February 5, 1999

Through, in effect, some double-blind experiments and lots of interviewing, we concluded that deploying this capability in the field organization in France reduced their parts cost five to ten percent, and reduced time spent at customers' sites five to ten percent as well. If you dare to extrapolate that to a worldwide sales organization with 20,000 people and go at the high end (ten percent), you save, in some way, 2,000 people. We have ten percent more productivity without adding people. Our people are now much more effective, faster, and get more work done. Also, anecdotal evidence shows that customers are much more satisfied. Interviews with employees say they like this. This, again, is a win-win situation for everyone.

We can see examples of that in other places. Organizations are using new technology to fix problems. That is exactly what happened in France. The field service organization went from being a problem to one of our best European field organizations. It is a really good story about how all this can come together.

There is a lot of research about the effectiveness of work groups and teams. We are not talking about a team here. Of course, all of France, or all 20,000 people, is not a team. What do we have then? We call it a community of practice. AAAS is a kind of community of practice—a group of people with common interests. They share knowledge and get

Figure 2
Community-Owned, Maintained, Knowledge Environments

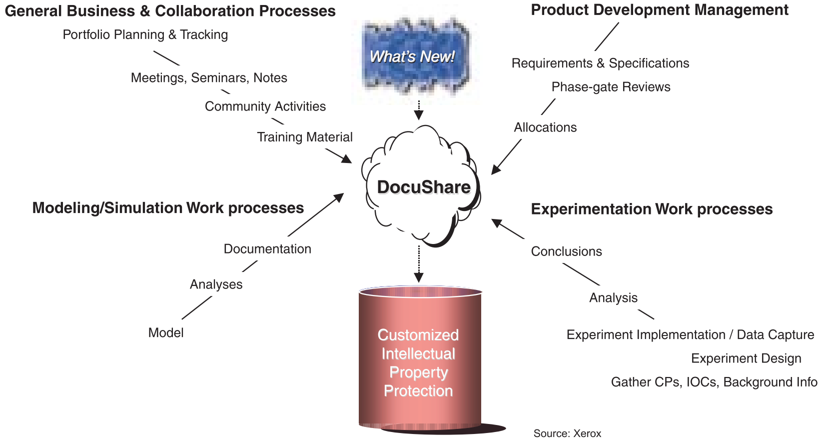

together in conferences and interact over the telephone. This accelerates the process and provides real-time capabilities to further the area of interest. But this is more of a notion of mass migration than it is a work group or team.

Knowledge management projects are as tricky and as difficult to do well as any other technology or social engineering project within an organization. One study looked at about 100 projects to compare high-impact projects with low-impact projects. Not surprisingly, a project strategy with an implementation plan is associated with successful projects. If you do not have such a strategy, it is just like anything else, you are probably not going to be successful.

In 100 percent of the high-impact projects the tool or the capability put in place tells the user what to do next. In low-impact projects the percentage is much lower. The conclusion is that projects are successful if you tell people who have some information what to do next. This is based on a lot of people's experience.

In the social aspect of knowledge sharing implementations, a very high percentage of high-impact projects are supported by changing management processes that help to implement the whole activity. Literally none of the low-impact projects included thinking on change management, that is, the people part of this whole effort.

To summarize, knowledge management is a major enabler for organization effectiveness. Many technical and social forces are at work and they are improving decisions and action. Decision making is being reshaped by low-cost technologies and change management techniques that change work practices. The combination of those two things is producing great results.

More information is available at http://www.mcse.external.xerox.com/docushare. Type "AAASguest" for user name. The password is also "AAASguest." Click on "knowledge."

15 Unnatural Acts: Building the Mature Firm's Capability for Breakthrough Innovation

Richard Leifer and Mark Rice

Background and Motivation

The contemporary competitive environment has been and continues to be driven by technological revolution, globalization, hypercompetition, extreme emphasis on price, quality, and customer satisfaction—with a resultant increasing focus on innovation (Hitt et al., 1998). These factors cause increased competitive complexity and dynamism, requiring an increasing emphasis on innovation as a strategic competence.

The competitive landscape began to change drastically during the 1980s, when technology-rich U.S. firms experienced increasing difficulty competing in various industries such as semiconductor memory chips, office and factory automation, and consumer electronics (Morone, 1993). Some have attributed these failures to an inferior capacity, relative to competitors, for achieving continuous and incremental improvements in products and processes targeted at increased quality and decreased costs. In response, in the past decade U.S. firms have focused on incremental innovation in existing products or processes, with an emphasis on cost competitiveness and quality improvements (Betz, 1993; Hamel and Prahalad, 1991; Morone, 1993).

Richard Leifer is an associate professor, and Mark Rice is an assistant professor at the Lally School of Management and Technology, Rensselaer Polytechnic Institute. *This article is based on remarks delivered at the 24th Annual AAAS Colloquium on Science and Technology Policy, held April 14–16, 1999, in Washington, DC.*

There are indications, however, that a focus on incremental innovation is an incomplete approach. Firms that hold the largest market share in one product generation often fail to maintain leadership when technologies shift (Bower and Christensen, 1995). Long-term competitive advantage comes from conflicting, yet ultimately complementary, activities such as maintaining a steady stream of incremental and continuous improvements in established business lines, and setting aside existing successful products for new innovations (Hitt et al., 1988). Thus, developing new businesses and products based on breakthrough (also termed discontinuous, radical, or game-changer) innovations becomes critical to long-term competitive advantage.

The distinction between incremental and breakthrough innovation probably stems from the definition of innovation proposed by Schumpeter (1934), who suggested that an innovation idea is the catalyst for a departure from existing practice supplanting standard operating procedures. March (1991), drawing on Schumpeterian themes, makes a distinction between exploration and exploitation. Exploitation has to do with refining or expanding existing products or processes, while exploration has to do with something fundamentally new, including new products, processes, or combinations of the two. Hence, breakthrough innovations are those that depart from the past and result in new products or services. Incremental innovations usually emphasize cost or feature improvements on existing products or services (Betz, 1993; Hamel and Prahalad, 1991; Morone, 1993). In contrast, breakthrough innovation concerns the development of new businesses or product lines based on new ideas or technologies (Morone, 1993) or substantial cost reductions that transform the economics of a business.

Although there may be significant risk and uncertainty associated with incremental innovation, generally uncertainty is more identifiable and manageable compared to breakthrough innovation. Hence, management practices utilized in the development of incremental innovations can be clearly defined and systematized. With respect to project planning, for example, a number of researchers have identified various stage-gate processes for incremental innovation development, which are widely understood and utilized in our sample companies (Cooper, 1990; Marquis, 1969). Since uncertainty is relatively proscribed, managers can define a continuum of project activities punctuated by a series of decision-making gates—where go/no-go decisions can be made.

By comparison, the uncertainty associated with breakthrough innovation is much greater. There is uncertainty about whether the technology will work, what the markets are, and what the applications might be. This uncertainty is difficult to manage in a systematic manner. Although the concept of breakthrough innovation frequently appears in the literature (Wheelwright and Clark, 1992) in terms of its importance or definition, there are relatively few ideas of how to manage it.

Block and MacMillan (1993) and McGrath and MacMillan (1995) point to differences between managing incremental and breakthrough innovation. It has been noted that breakthrough and incremental innovation processes require different organizational capabilities (Henderson and Clark, 1990). Approaches to market learning are fundamentally different as well (O'Connor and Veryzer, 1998). Hence, we contend that the processes and practices appropriate for managing incremental innovation may be inadequate, and, in some cases, dysfunctional for the development and commercialization of breakthrough innovation where uncertainty, risks, and potential rewards are much higher.

Attempting to migrate the disciplined processes developed for incremental innovation into the breakthrough innovation domain may be inappropriate due to differences of the phenomena being managed. Managers of the companies in our study reported difficulty in trying to do this. The lack of clear managerial prescriptions in the literature and the confusion of the managers in our sample companies of how to manage breakthrough innovation provide the motivation for this study. In this chapter we describe and categorize a variety of practices in use—and in so doing, uncover critical issues facing managers with indications of approaches to solving them. Our goal is to provide managers with a starting point for experimentation with new practices and tactics for improving breakthrough innovation effectiveness and to provide researchers with fruitful directions for further study.

In this study, we define a breakthrough innovation project as a formally established project with an explicit budget and organizational identity; and one that is viewed as offering the potential for a factor of 5–10 times (or more) improvement in product performance, an entirely new set of product performance features, or a significant (>30%) reduction in cost. This definition was derived from three sources: a review of the various definitions of breakthrough innovation in the literature (e.g. Leifer, 1997), discussions among the research team members based on the various perspectives of their disciplines, and intense discussion with industry representatives on the Discontinuous Innovation Subcommittee of the Research-on-Research Committee of the Industrial Research Institute.

Research Methods

The project was a multiple case, four-year longitudinal study of twelve ongoing breakthrough innovation projects (1995–1998) in ten large, mature firms: Air Products, Analog Devices, DuPont, GE, GM, IBM, Nortel, Otis Elevator, Polaroid, and Texas Instruments. The study was underwritten by the Sloan Foundation and sponsored by the Industrial Research Institute (IRI), a professional association of Fortune 500 research and development managers. The multiple case study design allowed us to explore the similarities and differences of management practices across industries, companies, and projects.

The study team included six faculty and several doctoral students from various disciplines. The study was deliberately designed as an exploratory, theory-building study rather than a conventional, theory-testing study. A panel of informants from each company met with the study team for in-depth interviews. The informants typically comprised the technical inventor or discoverer, project manager, project champion, and a senior manager. Each interview was taped and transcribed, resulting in many thousands of pages of transcripts. In addition, the research team had access to project documentation, reports, and business cases.

Multiple Comparison Case Study Methodology

This research project employs a multiple case study methodology. Case study research involves the examination of phenomena in natural settings. The case study method is especially appropriate for research in new topic areas, with a focus on "how" or "why" questions concerning a contemporary set of events (Eisenhardt, 1989). Case study research that employs multiple cases should follow a replication logic (Yin, 1994). The complexity of case study research and the high level of interpretation that is necessary create an advantage for the use of research teams. Multiple investigators can bring to the research a variety of experiences and complementary insights. A mix of different perspectives can increase the likelihood of discovering novel insights. Convergence of opinions from various researchers can enhance confidence in the findings, and conflicting views can keep the research from premature closure (Eisenhardt, 1989).

Field Study Sample Selection

Each participating firm nominated one or more projects for research consideration. To be included, a project was viewed by the research and development (R&D) and project managers of the firm as a "game-changer"; i.e., that it meets the characteristics of breakthrough innovation defined above. The project had to be formally established—with a project team and a budget. The participating companies in conjunction with the research team selected one or two projects for inclusion in the study. Although we gathered contextual information at the firm level, the project was the unit of analysis.

At the point we began gathering data, the projects were in various phases of development, although none were close to commercialization. As will become clear below, it is difficult to characterize the projects' stage of development because breakthrough projects do not seem to follow clear stages of development. Over the four-year study, the following projects were commercialized or are now in the public domain: BIOMAX (a DuPont degradable material); the TI digital mirror device; Analog Devices' air bag actuator; the IBM SiGe chip; and the Otis Elevator bi-directional elevator. For the remainder, the identities of the breakthrough innovations are cloaked to protect the competitive position of the companies.

Findings and Their Relationship to the Literature

We begin with an opening discussion of the nature of the breakthrough innovation life cycle to provide a context for our findings.

The Nature of the Breakthrough Innovation Life Cycle

Our team developed timelines for each project. The timeline for Project 7, presented in Figure 1, provides a typical picture of the general findings related to the breakthrough innovation life cycle derived from all twelve field projects.

In this case, the stream of technology R&D activities that preceded the formation of the commercialization project team began in 1980. The lines represent the applications pursued, and the thickness of the lines represents changes in the level of commitment of human and financial resources. The life cycle diagram indicates that Project 7 waxed and

Figure 1
Timeline for Project #7.

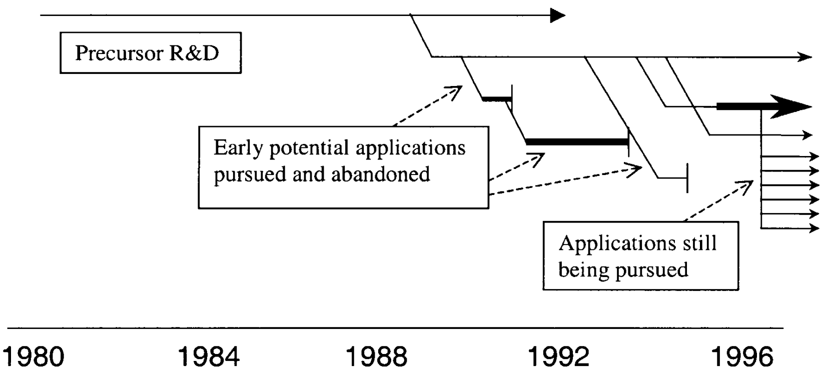

waned throughout its history. In the late 1980s and early 1990s, several streams of development activity reached dead ends and were abandoned, requiring the team to generate ideas for new high potential avenues to explore.

The development effort often took off in directions unanticipated at the project outset, particularly in response to discontinuous events related to changes in personnel, the formation of internal and external alliances, and key funding milestones. In late 1996, the technology was successfully transferred from the R&D business development organization to a business unit for manufacturing. Uncertainty about target markets and requirements for adoption by lead users remains high; hence, a variety of market applications continue to be pursued. Though initial target applications were not judged to be adequate for achieving an acceptable threshold of revenues, initial market response was sufficient to warrant moving forward. In fact, one of our follow-up interviews revealed that visibility in the marketplace derived from initial product launches stimulated the identification of new and more promising application domains suggested by potential customers.

This project is typical of all the projects studied. The life cycle of breakthrough projects is profoundly different from that of continuous improvement projects. They are long-term (typically ten years or longer) and sporadic, with many stops and starts, deaths, and revivals. Projects are non-linear. These projects were fraught with uncertainty about the technology, potential market applications, customers, and organizational/management issues. Finally, all the projects were context dependent. The development process was significantly influenced by personalities, personal preferences, organizational culture, and informal networks.

On the basis of this discussion, we conclude that management processes for breakthrough innovation are substantially different from those for incremental innovation. For example, we found that there is a desire to apply the stage gate system for managing product development developed by Cooper (1990). Cooper indicates that his system can be applied for the development of new products, not just extensions or incremental improvements. However, in practice, the stage gate system engenders a sense of linearity—with clearly specified stages and gates—whereas breakthrough innovation is characterized by non-linearity and uncertainty. The breakthrough implementation tactics we observed and which we describe below were designed to address the uncertainty and non-linearity problems of breakthrough innovations.

Tactics for Fostering Breakthrough Innovation

Our aim in this section is to examine our findings in the context of the strategy literature. We are focused on tactics that can be implemented in the organizations of today for increasing breakthrough innovation effectiveness. These five implementation tactics were based on an analysis of the empirical data. The first three tactics refer to necessary activities in the process of launching a radical innovation project, and the fourth tactic protects innovation projects from organizational pressure and resistance. The fifth tactic reflects our observations that, in current practice, breakthrough innovation relies on individual initiative rather than a systematic organizational process. The five tactics discussed are: stimulating attractive ideas; promoting opportunity recognition; evaluating and screening breakthrough innovations; creating incubating organizational structures; and catalyzing individual initiative.

Stimulating Attractive Ideas

Burgelman (1983) identified two kinds of strategic activities: those induced by the firm's current concept of corporate strategy; and autonomous strategic activities that fall outside the scope of the current strategy. It is useful to examine idea stimulation in our data set in the context of Burgelman's two kinds of strategic activities. The ten firms employed a variety of mechanisms to stimulate breakthrough innovation ideas that led, eventually, to the twelve projects. Among these were articulating holy grails, articulating strategic intent, issuing requests for proposals, promoting connections to external sources of technical information, conducting technology forecasting exercises, convening think tanks, creating an idea generation "sand box," and rotating talent.

The first three mechanisms were used by management to communicate to the entire organization the importance of idea generation. This first kind of strategic activity targeted at stimulating breakthrough innovation ideas can be thought of as a "strategic push." Six projects emerged partially due to strong articulation by senior management of the need for new ideas. The responses to these three mechanisms relate to Burgelman's (1983) first kind of strategic activity; i.e., activity induced by the firm's current concept of corporate strategy.

The remaining five mechanisms were used to enhance the idea-generating productivity of specific individuals or groups within the organization. They provide a context that supports an organizational response to "strategic push" as well as the second kind of strategic activities identified by Burgelman (1983); i.e., autonomous strategic activities that fall outside the scope of the current strategy. The specific ideas that are generated may fall within or outside of current corporate strategy, but the strategic choice by senior management to create a facilitative, supportive context for stimulating idea generation brings the activity itself into the strategic framework of the firm. We discuss below how several of these mechanisms were utilized.

Articulating Holy Grails. In seven of the twelve field study projects, there was a shared awareness among researchers and research managers of a technical "holy grail" within their industries. A "holy grail" is a potential technical breakthrough that is generally recognized as a catalyst for transforming an industry, if it can be achieved. For example, in the auto industry a dramatic improvement in fuel efficiency to 80 mpg would provide substantial competitive advantage to the firm that achieved the technical breakthroughs required for achieving this holy grail. In the people mover industry, current elevator technology limits the height of buildings. If that technical limitation could be overcome, it would be possible to erect buildings substantially higher than is possible today, such as the proverbial "mile high building." In the information storage industry, creating the capacity to have gigabytes of memory on a floppy disk was seen as a holy grail.

Articulating Strategic Intent. In four cases, articulation of strategic intent created a heightened awareness of the need for idea generation and alertness to new ideas. Although articulating strategic intent can take the form of articulating the importance of pursuing a holy grail as a strategic goal, it can alternatively be expressed in less specific terms—a more general call to arms. For example, in one case, the CEO pushed for new businesses in the white spaces between existing business units. In another, the CEO challenged the company to identify new applications of the firm's breakthrough technology for a market in which the company was not currently active.

Issuing Requests for Proposals. In five companies, competitive pressures threatening the core businesses of the firm or the perceived need for new strategic opportunities caused senior management to issue a request for proposals—either to the company or to a business unit—to pursue breakthrough innovations. In two cases, hundreds of proposals were submitted for evaluation and screening.

Promoting Connections to External Sources of Technical Information. The innovation literature argues consistently that external information gathering (Utterback and Brown, 1972; Tushman and Nadler, 1986; Gluck, 1988; Betz, 1993; Martin, 1984) creates a foundation for idea generation. In five of the twelve field study cases, proactively connecting to external sources of information had a significant impact on the generation of the initial idea. External sources included a scientific paper read by a senior research scientist; ideas submitted by external sources; a scientific paper on a secondary area of interest for a scientist; a brown bag lunch with a university researcher; and awareness of an innovation in an adjacent market.

General Observations About the Practice of Idea Generation. With one exception, idea generating mechanisms were applied sporadically and in an ad hoc fashion—rather than systematically, continuously, and strategically. When the idea generation activity was generally undertaken through individual initiative, and not incorporated into institutional processes, that idea generation mechanism ceased with the retirement or departure of the individuals. Results indicate an opportunity for senior managers to create a facilitating context and strategic push as effective stimuli for the generation of breakthrough ideas.

Promoting Opportunity Recognition

We found that opportunity recognition was typically an event that followed idea generation and triggered the ensuing process of evaluation and screening. The critical role of opportunity recognition in the process of developing and commercializing a breakthrough innovation emerges from a number of sources. To take advantage of shifts in technology, market, and competition, the opportunity first needs to be recognized and interpreted within the context of the firm's environment (Myers and Rosenbloom, 1993). Since companies sensitive to environmental shifts are best positioned for discontinuous change (Utterback, 1994), success comes from an ability to imagine markets that do not presently exist, and then to invest in their development ahead of the competition (Hamel and Prahalad, 1991). To accomplish this, companies need to develop the ability to sense, communicate, and appreciate the early signs of change (Gluck, 1988).

It is interesting that all the researchers cited above refer to the capacity of the company, or of management in general, for opportunity recognition. In contrast, we found opportunity recognition was based on individual initiative rather than a capacity or practice of the firm (Rice and Kelley, 1997). Kirzner (1973) defines the pure entrepreneur as a "decision-maker whose entire role arises out of his alertness to hitherto unnoticed [profit] opportunities," or in the words of Rumelt (1984), "sources of potential rents." In research, opportunity recognition related to breakthrough innovation was highly dependent on individual initiative and capacity rather than on routine practices and procedures of the firm. It was not a random act, but was generally reactive in nature; and it was unusual rather than proactive and routine.

In nine of the twelve projects, a low- to mid-level research manager performed the initial act of opportunity recognition. In two of the twelve projects, senior scientists, who had control over discretionary resources to support initial feasibility testing of real or potential innovations, recognized the opportunity. In essence, these two individuals played dual roles—as researcher and research manager, rather than as a researcher only. In only one of twelve cases was a senior technical manager responsible for the initial opportunity recognition, and in no cases did a senior corporate manager fulfill this function. The initiative of individuals set in motion activities that resulted in the establishment of a project to commercialize the breakthrough innovation.

In the majority of cases, the scientists who envisioned, worked toward and/or accomplished the breakthrough innovations had some idea (typically general and somewhat vague) of the application domain(s) for their innovations, but limited understanding of the market (Rice and Kelley, 1997). In eight of the twelve cases, the idea generator was not the opportunity recognizer. The scientists' research managers, who recognized the opportunities associated with their breakthrough innovations, had sufficient understanding of the potential market or application, which, when combined with their technical expertise, allowed them to recognize the business opportunities. The frequent interaction between the first line research managers and the scientists who were the idea generators made it likely that they would be in the best position to learn of the breakthrough ideas and make the connections to potential business opportunities. Unless someone close to the idea understands its implications, there is little chance it will reach the radar screen of higher level managers.

The leap in thinking required of these research managers to recognize an opportunity is reflected in the comments of one research manager:

> We didn't know much about market size at the time. The market is big. It is enormous. It would be a killer technology in this application domain based upon my physical understanding of what would be required in that industry.

Breakthrough innovation opportunity recognition appears not to occur in a formal way, but rather is part of the organization's informal network and organizational culture. The report of an initial opportunity recognition that led to the establishment of one project is illustrative. Two research managers for two scientists involved in the initial research independently contacted two individuals within the business development group about an upcoming technical review of the research. They suggested that the business development people come to hear an about interesting idea. One of the two was too busy too attend. The other stated:

> I get a notice every day for all the technical reviews going on here. I've got so many other things to do that I don't go unless somebody tells me to go. I was invited in by the scientist's research manager who said that this technology has the potential for making product X. I would not have normally gone to that review.

It is likely that the initiation of this project would have been delayed, or missed altogether, absent the occurrence of the opportunity recognition event.

In all cases, the desire to pursue an identified opportunity caused the research manager to reach out to other parts of the organization for support or resources. Given the high degree of technical and/or market uncertainty, the research manager sought confirmation of his perception of the opportunity. In addition, successful pursuit of the opportunity would over time require a significant commitment of the company's resources, which in turn would require decision-makers with authority to commit resources to develop the innovation. The research manager became the catalyst for leading others in the organization to recognize the opportunity.

Evaluating and Screening Discontinuous Innovations

There was evidence in our study of competing for the attention and financial support of senior management via corporate requests for proposals and proposals to venture boards, but this competition was not driving the evaluation and screening process. In fact, the concern of managers involved in our study was focused on how to evaluate and screen innovations much earlier in the process, before making proposals to venture boards typically becomes an issue.

A critical facet of managing the breakthrough innovation process is knowing that the pursuit is worth the risk. R&D managers recognize that the evaluation process for fundamental new lines of business differs significantly from that of extension projects. Yet our field studies offered limited evidence that there is a deliberate process or strategy for evaluating these projects differently. Screening was either undertaken as part of the normal project evaluation process, or was treated in an ad hoc fashion. We note that in some cases, traditional evaluation criteria and methods were used, but were generally not perceived by the technologists and managers at the project level to be relevant to them.

The primary purpose of screening and evaluation activities was to support decision-making related to resource (financial and human) allocation and, as necessary, acquisition. In all twelve cases, a positive outcome of the initial screening and evaluation process caused a commitment of internal human and financial resources from the discretionary funds of a research manager to continue exploratory work. At this point, uncertainty on one or more dimensions (technical, market, organizational, financial) was so high that the initial commitment simply reflected a conviction on the part of the research manager that the magnitude of the potential opportunity was large enough to warrant additional technical development.

In all cases, evaluation and screening at the project level was ongoing, initially centering on achievement of technical milestones, failure to achieve those milestones, or discovery of unanticipated technical hurdles (O' Connor and Veryzer, 1998). It also occurred with respect to market learning that came through interactions with early adopters and through the experience of lead users with prototypes. By comparison, evaluation and screening by senior management occurred in the context of periodic budgetary reviews and was conducted by an individual senior manager, a project "board of directors," and/or a venture review board, but still based on primarily technical considerations rather than on formal "business cases."

For breakthrough innovations, initial assumptions could be made about market- and economics-related issues, but generally there were one or more critical issues related to technical uncertainty that took precedence. These needed to be resolved in order to be able to embody the technology into some sort of prototype that could be used for market learning. Hence, the focus of screening and evaluation activities in early project stages was predominantly technical. The primary evaluative questions driving the radical innovation projects in this study included the following:

- What is the magnitude of the impact this technology can have on the market?
- What will this technology enable?
- Can this technology deliver the magnitude of benefit that is needed?
- What are the technical hurdles we must overcome to get this thing (or process) to work?
- What are the projected performance characteristics?
- What yield from the manufacturing process must be achieved to make the economics of the business attractive?

The focus of this set of questions was on the return of new value, in a variety of ways, to the market (O' Connor and Veryzer, 1998). The long-term profit potential was assumed to be significant even though not currently quantifiable. In comparison, typical screening criteria for an incremental improvement investment is how much promise the project offered to the firm, characterized as, "What is the profit impact?" or "How fast will it grow?" or "How much market share can we expect to grab?" The focus of each of these questions is on a return to the firm over a given (usually specified) and relatively short time frame.

For breakthrough innovations we found an emphasis on a more experimental, hands-on approach and a reliance on past experience to assess the value of the technology to the market (Rice et al., 1998). Potential customers, or "lead users," were not the only vehicle for this. Perspectives of many constituents were sought, including leading members of the technical community, senior management at both the corporate and business unit levels, and line managers connected to the current customer base. There was a heavy reliance on "probes" to potential early adopters and others who were relatively sophisticated in relevant technical arenas, including: professional conferences and meetings, where data are presented for the technology community's reaction and to gain potential customer interest; the demonstration of the product via early prototypes for reaction within the firm; and evaluations by potential customers of early working versions over extended trial periods.

These probes were more experimental in nature than analytical, and they were designed for technical and market learning more than market evaluation (O' Connor and Veryzer, 1998). The purpose was not to assess the impact on sales, but rather to assess the degree to which potential users will experience value-in-use. Typically, several potential applications were pursued, usually serially, to test technical and market assumptions. Positive results of these activities were critically important for gaining support internally, both to sustain financial support of senior management and to increase the receptivity of business unit managers, who would eventually manage the product emerging from the project.

Creating Incubating Organizational Structures

Galbraith (1983) and Quinn (1985) state that organizations pursuing innovation require a specially designed structure that enables them to develop significant innovations not consistent with the existing organization concept. The appropriate structure is likely different in various situations and at different times. A structural design may need to be changed as conditions change (Twiss, 1986). A number of authors recommend that separate structures should be maintained: an operating structure for routine lines and an innovating structure for new innovative products (Galbraith, 1982; Burgelman and Sayles, 1986; Kanter, 1989). They argue that the mainstream mode of operation is too slow and conservative to allow for development of projects characterized by high uncertainty, high intensity, and high autonomy. Innovative new projects are not likely to benefit from the existing organization's experience base (Kanter, 1989).

Breakthrough innovation seems to work best, especially in the front end of the process, when separated from ongoing business activities. For most operating businesses, breakthrough innovation is an unnatural act because the uncertainty is too high, the time horizon too long, and the investment too large given the inherent risks. Breakthrough innovation projects are badly aligned with the reward structure of operating businesses. The costs occur in the present and benefits do not arrive for perhaps ten years or more. Thus, regardless of potential long-term benefit for the firm, within the short time horizons of operating units the impact of breakthrough innovation is negative; i.e., depressing short-term profitability. For all these reasons mechanisms for protecting projects—incubating arrangements—functioned to allow projects to develop outside of organizational pressures and to find resources not normally associated with project development.

We found that successful breakthrough innovations developed in a variety of incubating arrangements. These incubating homes allowed innovations to develop enough maturity to be attractive to the operating units while being protected from short-term organizational performance metric requirements. The relative effectiveness of the various options varied with context. In early development, the home of the incubating organization is typically corporate research and development (CRD), a new business development group operating within CRD, a business unit, or a new ventures group operating at the corporate level.

Incubating arrangements frequently operated across internal and external organizational boundaries to bring needed resources into the project. Partnering was used for risk reduction throughout the breakthrough innovation development process. Participation of internal and external partnering organizations varied but had a significant impact. A wide variety of partners were observed, including government. In eight out of twelve cases, government agencies were a major source of funds after the project was formalized. Government funds were used to extend, expand, and accelerate projects, but in only one case was government funding a trigger or motivation for the project.

There was widespread use of alliances for a variety of purposes, including manufacturing, application development, market probing, and joint development of technology (Rice et al. 1998). A broad spectrum of partners were involved, including other large firms, universities, government laboratories, and small high tech firms. Strategic alliances served to contribute knowledge about markets and technology as well as to help managers gain visibility and legitimacy for their projects. They aided in monitoring the environment, and provided access to related cutting-edge technologies that could provide entrance into new markets.

All of this points to the importance of a resource view of the breakthrough innovation process. Since these innovations typically require expertise, funding, and the use of organizational resources (such as time in a fabrication facility or on a manufacturing line) outside of normal resource allocation processes, the success of breakthrough innovations depends on the ability of project managers, champions, and sponsors to find and acquire those resources.

Catalyzing Individual Initiative

In contrast to the literature that describes institutionalization of the incremental innovation process, we found that successful breakthrough innovation depended on the actions of individuals. Rather than relying on organizational systems to manage the process (as might be expected with incremental innovation), individuals were the prime movers and sustainers. We identified several types of key actors—creative scientists, opportunity recognizers, multiple champions and supporters, and project team leaders and members. All of these individuals demonstrate a common characteristic—a passion for and belief in the innovation. This passion overcomes and protects the innovation from organizational forces that naturally resist new ideas.

Creative Scientists or Engineers. Competitive advantage through technical leadership requires recruiting and development a cadre of talented scientists and engineers who have both technical skills and creative, out-of-the-box, idea-generating capabilities (Morone, 1993). In nine of the twelve projects, a single hero scientist or creative engineer had the initial technical insight that set off the chain reaction of events leading to the breakthrough innovation. Often, it was a creative, cognitive act that linked disparate bits of information together.

In organizations that had positions of research fellows, hero scientists (technical champions) emerged from the ranks of senior research fellows or technical people with high prestige. The support and championing of the breakthrough innovation arena was tolerated and protected by their status. Nevertheless, there was still immense difficulty, and the majority of these hero scientists simply stuck out their necks when advocating their breakthrough innovations. This, of course, comes at a risk, but the champions we observed were so passionate about their projects that they paid scant attention to potential downside to their careers.

Opportunity Recognizers. Taking the technical idea and recognizing the potential application or market opportunity was a critical event within the breakthrough innovation process. As indicated earlier in this paper, in eight of our twelve cases first line technical managers (themselves experienced engineers or scientists), not senior managers, were the individuals who first recognized the opportunity, even though they were not the ones who came up with the idea. These innovation recognizers were convinced of the significant potential impact of the innovation and had the confidence to champion it in the larger organization.

We also found that first line managers were not always heard and sometimes gave up; i.e., the opportunity recognizers were not necessarily the champions. They may make the links and communicate the nature of the opportunity to someone else, but then choose not to fight the battles that so often arise with unfamiliar, uncomfortable, or alien topics. Thus, although a creative capacity was necessary, it was not sufficient to achieve progress. Also important was a determination to get the attention of the larger organization and to fight for the resources to pursue the opportunity.

Multiple Champions and Supporters. We observed the importance of multiple champions playing multiple roles. As Day (1994) and Venkataraman et al. (1992) found, champions play multiple roles or there are multiple champions playing different roles. As the literature suggests and our data confirm, champions played a key role in driving these projects forward, especially in the face of obstructionism from other parts of the organization or when intensity and perseverance were required to overcome hurdles.

We identified several types of champions: technical champions, project champions, senior management champions, and business unit champions. Multiple champions are necessary to support and protect breakthrough innovations because organizational resistance comes from all sides, and the projects by and large do not have organizational legitimacy (Dougherty and Heller, 1994). Although it might seem a reasonable strategy to somehow coordinate the activities of these multiple champions, we found little of this happening. Champions seemed to act individually, though simultaneously, in support of the innovations. While occasionally a single person assumed multiple championing roles, more often than not different people played each championing role. Finding champions was critical to breakthrough innovation success.

Informal networks emerged as critical in all twelve innovations for gaining resources. These networks occurred vertically and horizontally within the organization, including within R&D; between R&D and the business units; and between the company and customers, suppliers and government. Informal networking played a prominent role in idea generation; idea evaluation; generating political/financial support; gaining access to scarce resources, e.g. manufacturing capacity; connecting with friendly customers for alpha testing; and attracting government funding. Due to the innovation's need for legitimacy, protection and access to resources as discussed above, champions tied into organizational networks were critical to overcoming these barriers. One project champion captured the importance of informal networks. *"It's important to know who the right people are in the company to get assistance and support; there's a secret way the company operates. Because of the way in which managers have grown up around here, you have these internal networks and you work them and that's what makes the place work."*

Conclusions

Based on a four-year study of twelve breakthrough innovation projects in ten large, mature firms, we concluded there were significant differences in characteristics and management practices of breakthrough vs. incremental innovation projects. Best practice management of incremental projects suggests a well-defined, linear path with clear go/no go decision points; e.g., the stage gate model (Cooper 1990). In contrast, breakthrough projects are highly uncertain and unpredictable, non-linear (with lots of detours, starts, and stops), stochastic, and managed more through individual initiative rather than through formal, established organizational processes. Due to these dynamics, championing/spon-

Table 1
Summary of Implementation Tactics and Their Effects

Implementation Tactic	Implementation Tactic Effect
1. Stimulating Attractive Ideas	Ensures a stream of high quality ideas for triggering breakthrough innovations; create a strategic context for push, develop goals as pull.
2. Promoting Opportunity Recognition	Sensing, recognizing, and supporting potential opportunities; low- to mid-level research managers most likely to play this role as they are closest to ongoing research.
3. Evaluating and Screening Opportunities	Criteria for continuing and supporting project; early development evaluated on technical considerations rather than on traditional business cases.
4. Creating Incubating Structures	Protects the innovation from organizational resistance until it is strong enough to compete for resources and support on its own merits.
5. Catalyzing Individual Initiative	Technically creative individuals, champions, and supporters necessary for ideas, recognition, securing resources, and organizational support and protection. Getting innovation on radar screen of organizational decision-makers critical to continued project existence.

sorship of senior management and reduced reliance on financial and market numbers as criteria for project continuation (at least early on) seem critical for long-term project success. The front ends of these projects seem to be extended compared to incremental projects, with extensive exploring and experimenting, probing, and learning rather than targeting and developing. Our data further indicates that when projects reach a dead end in one line of exploration, they frequently cycle back to the front end again, exploring another idea or market opportunity.

We identified a set of implementation tactics for coping with these phenomena, as summarized in Table 1.

These implementation tactics were not applied in an isolated or sequential manner, but in a highly interactive manner. For example, both incubating organizational structures and senior manager championing reinforce each other, serving to protect radical innovation projects from the pressures and practices of the ongoing business. Idea generation is often the trigger for subsequent opportunity recognition—either by the idea generator or a second individual in close proximity—and opportunity recognition, in turn, stimulates a set of initial evaluations. Because of overall uncertainty in the early years of these projects, idea generation is an ongoing and interactive part of early opportunity recognition and evaluation/screening and development activities. Hence, these tactics are used to promote and protect the discovery process—which, in the short term, increases uncertainty—but that for successful projects will, in the long term, reduce uncertainty sufficiently so that the project can be transferred into an operating unit that utilizes traditional management practices.

Conventional management techniques are unsuitable until uncertainty is sufficiently reduced. It appears that the primary imperative driving these projects was to reduce uncertainty to the point where conventional management practices could be applied or where it became apparent the project should be abandoned. From our observations of the projects included in this study, it appeared that much of the breakthrough innovation process was not deliberately managed. While it may be true that conventional management techniques are inappropriate, we believe a more systematic approach can and will be developed. This more systematic approach, which can be derived from an understanding of the dynamics and characteristics of the breakthrough innovation project life cycle, can incorporate the five implementation tactics identified in this study:

- Senior management can develop and articulate the need for breakthrough ideas, delineate white space opportunities, and set "holy grail" challenges.

- Senior management can support opportunity recognition of breakthrough ideas with the recruitment, development, and placement of technically prestigious individuals, well networked in the company, to act as sentries and scouts identifying and ferreting out new ideas.

- Senior management can evaluate and screen breakthrough projects with different metrics compared with those used for incremental projects. Evaluation based on technical feasibility, at least early on, will keep a project alive against ongoing financial pressures. Further, senior management can support alternative methods for accessing information used in the evaluation and screening process; e.g., experimentation and probes with potential customers for early market data.

- Senior management can support and protect breakthrough projects by developing incubating organizational structures in which projects can mature sufficiently to be able to withstand the scrutiny of more traditional project evaluation reviews. These incubating arrangements may be implemented within R&D centers, business development groups, or separately run departments or divisions. Finding partners and creating alliances with complementary technologies or capabilities supports resource acquisition and helps spread the risk inherent in breakthrough projects.

- Individuals play a critical role in developing and sustaining breakthrough innovations. Even in companies with effective incubating arrangements, senior management needs to support and protect breakthrough innovation projects visibly and vocally. Senior management can cultivate a set of individuals who operate within the organizational system and culture but at the same time challenge the organization by questioning current products and processes, thereby providing a catalyst for the initiation of breakthrough innovations and a medium for their development. Senior management can assist the process by making sure these unconventional (sometimes difficult) individuals are identified, supported, and rewarded. However, unless senior management plays the roles defined above, these out-of-the-box thinkers and actors may not emerge to play their roles.

Breakthrough innovation activity by its very nature is largely incompatible with mainstream business activity. Without stimulation and sponsorship by senior management, the flow of breakthrough innovation activity will be severely restricted by the culture and the drivers that serve to optimize mainstream business activity. Without the protection of senior management via championing and the creation of incubating structures and mechanisms, breakthrough innovation will fail to reach maturity.

Problems of organizational legitimacy result in difficulties in obtaining necessary resources, funding, personnel, time, support, access to organizational support services, and the like. In a sense, the success of breakthrough innovations depends on obtaining a minimum support level of these resources. Most of the mechanisms for supporting the project—multiple champions, informal networks, hero scientists, top management support, discretionary resources, and external alliances as well as government funding—are designed to obtain needed project resources. Many project members agreed that without government resources, for example, the project would have died of malnutrition.

There are researchers who are skeptical about the capability of large established firms to pursue breakthrough innovation. Our findings suggest that senior management has the opportunity to exercise strategic decision-making that can result in fostering, tolerating, or squelching the development and commercialization of breakthrough innovations. The five implementation tactics identified in this study can be used to proactively manage the balance of core, mainstream business activity vs. the breakthrough innovation activity that can provide the seeds for future business development.

Our study extends the existing literature by providing a foundation for developing appropriate management practices, rather than simply highlighting their existence and importance. We hope that this research effort will stimulate the development of additional managerial techniques and a systematic approach for implementing them that will allow firms to maximize the yield from their breakthrough innovation initiatives.

References

Betz, Frederick. *Strategic Technology Management*, McGraw-Hill, New York, 1993.

Block, Zenas and MacMillan, Ian C. Corporate Venturing, Harvard Business School Press, Boston, 1993.

Bower, Joseph L. and Christensen, Clayton M. Disruptive technologies: Catching the wave, *Harvard Business Review*, January-February, 43–53; 1995.

Burgelman, Robert A. A process model of internal corporate venturing in the diversified major firm, *Administrative Science Quarterly*, 28(2), 223–244; 1983.

Burgelman, Robert A. and Sayles, Leonard, R. *Inside Corporate Innovation*, The Free Press, New York, 1986.

Cooper, Robert. The new product process: A decision guide for management, *IEEE Engineering Management Review*, June, 19–33; 1990.

Day, Diana L. Raising breakthroughs, *Organizational Science* 5(2); 1994.

Dougherty, Deborah and Heller, Trudy. The illegitimacy of successful product innovations in established firms, *Organizational Science* 5(2): 200–218; 1994.

Eisenhardt, Kathleen M. Building theories from case study research, *Academy of Management Review* 14(4): 532–550; 1989.

Galbraith, Jay R. Designing the innovating organization, *Organizational Dynamics* 10(3): 4–25; 1982.

Gluck, Frederick W. Breakthrough innovation through creative leadership, in R. L. Kuhn, *Handbook for Creative and Innovative Managers*, McGraw-Hill, New York, 1988.

Hamel, Gary, and Prahalad, C.K. Corporate imagination and expeditionary marketing, *Harvard Business Review* 69(4): 81–92; 1991.

Henderson, Rebecca M. and Clark, Kim B, Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms, *Administrative Science Quarterly* 35: 9-30.

Hitt, Michael, Keats, Barbara W., and DeMarie, Samuel M. Navigating in the new competitive landscape: Building strategic flexibility and competitive advantage in the 21st century, *Academy of Management Executive*, 12(4): 22-42; 1998.

Kanter, Rosabeth Moss. Swimming in new streams: Mastering innovation dilemmas, *California Management Review*, 31(4): 45-69; 1989.

Kirzner, Israel M. *Competition and Entrepreneurship*, The University of Chicago Press, Chicago and London, 1973.

Leifer, R. Organizational and managerial correlates of radical technological innovation, in Dundar Kocaoglu and Timothy Anderson (eds), *Innovation in Technology Management*, Portland International Conference on Management of Engineering and Technology, 134–137; 1997.

March, James. Exploration and exploitation in organizational learning, *Organizational Science*, 2(1): 71–87; 1991.

Marquis, Donald G. The anatomy of successful innovations, *Innovation*, November 1969.

Martin, Michael J. C. *Managing Technological Innovation and Entrepreneurship*, Reston Publishing Company, Inc., Reston, VA, 1984.

McGrath, Rita Gunther and MacMillan, Ian C. Discovery Driven Planning, *Harvard Business Review*, July–August, 1995.

Morone, Joseph G. *Winning in High Tech Markets*, Harvard Business School Press, Boston, 1993.

Myers, Mark B., and Rosenbloom, Richard S. Research management and corporate renewal, *Conference on the Future of Industrial Research*, February 1993, Harvard Business School.

O'Connor, Gina Colarelli and Veryzer, Robert W. Linking advanced technologies to market opportunities: The nature and process of market visioning, Product Development Management Association 1998 Conference Proceedings, 67-78; 1998.

Quinn, James Brian. Managing innovation: Controlled chaos. *Harvard Business Review* 63(3): 73–84; 1985.

Rice, Mark P. and Kelley, Donna J. Opportunity recognition in the domain of discontinuous innovation: An exploratory study, presented at the 1997 Babson College-Kauffman Foundation Entrepreneurship Research Conference, 1997.

Rice, Mark P., O'Connor, Gina Colarelli, and Peters, Lois J. Managing discontinuous innovation, *Research Technology Management* 41(3); 1998.

Rumelt, Richard P. Toward a strategic theory of the firm, in *Competitive Strategic Management*, Prentice-Hall, Englewood Cliffs, New Jersey, 1984.

Schumpeter, J.A. *The Theory of Economic Development*, Harvard University Press, Cambridge, MA, 1934.

Tushman, Michael L., and Nadler, David Organizing for innovation, *California Management Review* 28(3): 74–92; 1986.

Twiss, Brian C. *Managing Technological Innovation*, Pitman Publishing, London, 1986.

Utterback, James M. Breakthrough innovation and corporate regeneration. *Research Technology Management* 37(4): 10; 1994.

Utterback, James M., and Brown, James W. Monitoring for technological opportunities, *Business Horizons* 15(5): 5-15; 1972.

Venkataraman, S., MacMillan, Ian C. and McGrath, Rita Gunther. Progress in research on corporate venturing, in D.L. Sexton and J.D. Kasarda (eds), State of the Art of Entrepreneurship, PWS-Kent Publishing Co., Boston, 487-519; 1992.

Wheelwright, Steven C. and Clark, Kim B. Creating project plans to focus product development, *Harvard Business Review* Mar–Apr; 1992.

Yin, Robert K. *Case Study Research*, Sage Publications, Thousand Oaks, CA, 1994.