PERGAMON

Evaluation and Program Planning 22 (1999) 91–103

EVALUATION
and PROGRAM PLANNING

The case study as research heuristic: lessons from the R&D value mapping project[1]

Barry Bozeman, Hans K. Klein*

School of Public Policy, State Data and Research Center, Georgia Institute of Technology, Atlanta, GA 30332, U.S.A.

Accepted 1 August 1998

Abstract

How can case studies be used as a research heuristic? If prototype case studies are performed, what can researchers expect to learn from them and how can they be structured to enhance their learning value? This paper considers that question and the learning from two case studies intended to inform multiple case studies undertaken later in the project. Two prototype cases are presented, one Brookhaven National Laboratory, the other from Los Alamos National Laboratory, each having as its objective providing information about how to design and execute the subsequent 30 case studies to be undertaken. This paper summarizes the cases, presents some of the lessons learned for the subsequent larger project and then considers more generally the use of prototype case studies and the preconditions for their successful deployment. Prototype case studies are particularly useful for helping set boundaries for later studies, identifying the ways in which the research setting affects research findings, making judgments about the accessibility and availability of data, and determining respondents' reactions to the research and the researchers.

Keywords: Case study; Methodology; Research evaluation; R&D value mapping; National laboratories

1. Introduction

Often policy evaluators and other social scientists approach case studies with considerable wariness. Even the titles of papers dealing with case study methods reflect ambivalence. One paper refers to case studies and related qualitative approaches as an 'attractive nuisance'; another considers 'Case Studies as a Serious Research Strategy' (Yin, 1981), as though case researchers were in need of a reminder to be less casual. This defensive posture pervades the content, not just the titles, of case studies and primers on case method. While more than one issue accounts for this defensiveness, perhaps the single most important one is the widespread suspicion that case study methods are insufficiently general and theoretical, that case studies are good for explaining the unique, but not up to the task of providing generalizable explanations (Dyer & Wilkins, 1991).

Practicing case researchers typically voice much greater confidence about the rigor of case studies and their potential for generating and testing theory (Eisenhardt, 1989; Lee, 1989a; Kennedy, 1976; Glaser & Strauss, 1967). Nevertheless, even optimistic case researchers recognize the difficulties in using case studies for theory development. Eisenhardt (1989:547) notes one of the more severe problems, that "a hallmark of good theory is parsimony, but given the typically staggering volume of rich data, there is a temptation to build theory which tries to capture everything."

If the 'staggering volume of rich data' is an occupational hazard among case researchers, most of whom work with one or a handful of cases, one easily comprehends the depths of the problem when dealing with scores of cases. In the 'R&D value mapping project' (Bozeman & Kingsley, 1997), the analytical approach *requires* more than thirty cases. In the latest phase of the project, some 50 cases studies will be developed (Bozeman, 1997) by teams of case researchers. In working

* Corresponding author. Tel.: +1-404-894-2258; fax: +1-404-894-0535; e-mail: hans.klein@pubpolicy.gatech.edu

[1] This paper is part of the R&D Value Mapping Project funded by the Department of Energy, Office of Basic Energy Sciences, contract no. 04371-23. The views presented here are solely the authors' and do not necessarily reflect those of the Department of Energy or Georgia Institute of Technology. The authors are grateful to Dave Roessner, Juan Rogers, Francisco Donez and Gordon Kingsley for their contributions to the larger project and to this paper. Lynn Austin provided assistance with organizing the case study literature. Requests for full case studies and other research products from the R&D Value Mapping Project should be sent to: RVM Project, State Data and Research Center, Georgia Tech, GCATT Building, 250 14th Street, Atlanta, GA 30318-0490, U.S.A.

0149–7189/99/$ - see front matter
PII: S0149–7189(98)00045–7

with so many cases, and with multiple researchers, the embarrassment of riches problem can be overcome only with analytical discipline and, especially, development of theory and hypotheses to guide the case studies.

An advantage of performing a large number of cases is the possibility of using early ones as prototypes. The function of case studies for theory development is widely recognized. Whereas the use of case studies for testing theory remains controversial, there is little dispute about their value for developing theory. In most instances where case studies are used to develop hypotheses and theory, they are used in anticipation of applying some other method such as survey analysis of field experimentation. But in the RVM project, early case studies have been used to develop theory and subsequent methods for larger numbers of case studies performed later in the project.

The chief focus of this paper is on the question 'How can case studies be used as a research heuristic?' That is, if one or a few early, prototype case studies are performed, what can we expect to learn from them and how can they be structured to enhance their heuristic value? Before turning to those questions and answers provided by RVM prototype cases, we consider some of the broader questions in the research uses of case studies, focusing particularly on their uses in our domain of concern, R&D impacts.

2. The case study tradition in R&D impact evaluation[2]

Yin has summarized the major strengths and limitations inherent in all case study designs. Since his treatment and other general methodological critiques of case studies are familiar, we begin with just a few general observations. First, case studies are useful for addressing questions regarding how and why phenomena behave, often providing hypotheses about behavior rather than validating general claims about behavior. Case study research often reveals a rich detail of information that highlights the critical contingencies that exist among the variables. The method is especially useful for exploration of topics when there is not a strong theory to which one can appeal.

The problems with case studies pertain to lack of rigor or, more to the point, a different meaning of rigor. The case study (not unlike most approaches to social research) allows equivocal evidence or biased views to influence the directions of the findings and conclusions. Related, case studies often extend little hope of valid external generalization. (However, Yin argues that concerns regarding the generalizability of case studies are exaggerated and outlines ways to remedy this criticism.) A third concern is that it takes a great deal of time and money to collect and analyze case study data, especially when case researchers are serious about addressing the problems of validity and reliability. Related, it also is an expensive method to conduct. The cost of the case study method reduces its utility for certain research problems.

The use of case studies for evaluations of R&D impacts has been shaped by two research questions: (1) What are the linkages between R&D and economic growth? (2) Are R&D projects meeting the policy objectives established for the organization that mandate addressing linkages between R&D and the economy? Answers to the first question have been the preoccupation of policy-makers since World War II when the impact of science on the welfare of the nation became dramatically clear. Answers to the second question have been the preoccupation of industry and government agencies who must demonstrate the economic benefits of specific R&D projects.

Initially, the case study method was employed to develop concepts and methods that would allow a more precise understanding of terms such as invention, innovation, technology transfer, or basic, applied, and development research. The ultimate thrust of this research was to develop concepts and methods that would allow a more explicit and thoughtful articulation of the causal relationships that link R&D and the economy.

There have been four genre of case studies used in the post-World War II era for examining the impacts of R&D. Three are different forms of retrospective analysis: (1) historical descriptions; (2) research events studied; and (3) matched comparisons. The fourth is a combination of retrospective analysis with other methodologies such as aggregate statistics, peer review, bibliometrics, and econometrics. Though the development of these case study types are roughly sequential and build upon the limitations of earlier studies, the development of new techniques has not resulted in the obsolescence and retirement of the earlier approaches.

The earliest approach was to conduct historical descriptions of the development of a specific technology. The work of Jewkes, Sawers and Stillerman (1969) is an example of this genre of case study, examining the relationship between R&D and innovation by tracing innovations back to fundamental supporting inventions. Similarly, Carter and Williams (1957) examined the stages in the generation and application of scientific knowledge from basic research to the commercial decision of innovation investment. Though historically informative, this approach did not result in a structured analytic framework with well defined concepts and methods of measurement.

From the 1960–70s, a series of massive case study projects was sponsored by government agencies in an effort to understand the linkages between R&D and economic growth. Studies such as Project Hindsight sponsored by the Department of Defense, and the Technology in Retrospect and Critical Events in Science project

[2] This discussion draws from Bozeman and Kingsley (1997).

(TRACES) sponsored by the National Science Foundation, further developed the analytic techniques used in retrospective analysis by identifying 'research events' in the development of specific technologies. Research events are defined as the occurrence of a novel idea and subsequent period where the idea is explored. Thus, the technique was to take specific research technologies and divide them into the research events that led to the successful development of the technology. Another development in retrospective analysis was to compare innovations that had been determined a priori to be of different types. For example, Project SAPPHO conducted pairwise comparisons of innovations that were successes and failures in terms of commercial diffusion.

The empirical results of these studies dramatically conflicted, reflecting the interests of the organizations that had sponsored the studies. Nor did these studies establish a strong conceptual base upon which further research could build. Economic and bibliometric techniques began to replace retrospective case studies as the preferred methods to examine the link between R&D and innovation.

Throughout this period, the case study method had also been used to evaluate the performance of specific R&D projects within the context of a policy objective. These objectives normally have an implicit, or explicit, assumption that R&D directly affects the economy. But the goal of the case study emphasizes evaluation of project performance in preference to developing contributions to theory. Though case studies seeking to establish linkages between R&D and the economy failed to establish a strong theoretical base, they nonetheless had a significant methodological influence upon case studies emphasizing project evaluations. Evaluation studies have mimicked the retrospective case study designs used to develop theory. For example, a recent case study conducted by Oak Ridge National Laboratory uses a form of retrospective analysis, charting the Department of Energy's contribution to the development of several building innovations.

The frustration with the findings from the case study design also led to a variation in case study research that combines several methodological techniques. As noted above, these multi-method approaches bring together case studies with peer review, bibliometric techniques, and econometric analysis under the heading of impact analysis. The goal of impact analysis is to look for levels of agreement between the different techniques employed.

3. Prototype cases studies and the RVM project

The research value mapping (RVM) approach was developed as a means of addressing some of the methodological weaknesses of case studies while, at the same time, preserving the richness and sense of context associated with case studies. By its very nature, RVM is resource-intensive and requires many interrelated cases. In light of the resources required and the difficulty of planning and coordinating multiple cases undertaken by multiple actors, RVM, even more than most research designs, benefits from planning and from learning during early phases of the research.

The use of prototype case studies is not uncommon for informing studies entailing other (non-case) methods. Thus, in the National Comparative R&D case studies of R&D laboratories (Crow & Bozeman, 1987; Bozeman & Fellows, 1988) were instrumental in developing questionnaires used in subsequent research (for an overview see Bozeman & Crow, forthcoming). Prototype case studies are less common when the subsequent method is additional case studies. The reason, of course, is that use of more than a handful of case studies is uncommon and, thus, researchers typically do not have the luxury of prototype cases.

4. The RVM project

While this paper is not chiefly about the RVM project and its research, but rather about the many uses of prototype cases, some knowledge of the project is relevant to understanding the narrower objectives of the prototype cases.

The RVM project's chief objective is to develop and apply valid approaches to assessing the economic and social impacts of research, including basic research. The project entails development and application of R&D value mapping (Bozeman & Roessner, 1995; Kingsley & Bozeman, 1997; Bozeman & Kingsley, 1997) an approach to combining qualitative and quantitative analysis of case data.

The RVM method yields an inventory of marginal benefits from R&D projects and empirical generalizations of the determinants of those benefits. R&D value mapping has much in common with earlier case study-based attempts to assess research, but is in many respects a significant departure. As in previous case studies of R&D impacts, RVM focuses intensely on particular projects and the events surrounding them. Case studies 'tell a story' about the chronology and events contained within the boundaries of the project and RVM is similar to traditional case studies in that it yields such a narrative. There is also an expectation that the case studies can contain a richness that goes beyond traditional aggregated quantitative studies to provide insights from detail and nuance. But RVM seeks to avoid some of the pitfalls of traditional cases.

Case studies are faulted as interesting stories which provide little systematic explanation. The RVM approach, beginning with carefully specified and testable models of causation, as well as a scheme for linking the cases, yields both particularistic and generalizable data. The particularistic information is much like that derived from a traditional case. The generalizable data comes from the quantification of elements across cases. RVM is similar to other case survey techniques whereby individual cases are scored by multiple coders, where resulting scores are subjected to inter-coder reliability analysis (Bullock & Tubbs, 1987; Larsson, 1993). Case scores are then categorized for pattern-matching both within groups of cases and between case groupings. Thus, each project 'tells a story' and, at the same time, gives rise to systematically measured observations. The research procedures of RVM are summarized in Table 1.

4.1. Prototype cases

Before undertaking the considerable task of conducting at least 30 intensive case studies, it seemed useful to the RVM research team to conduct *prototype* case studies to learn about the feasibility of the research plan and to refine case research methodologies and instruments. After summarizing two prototype case studies (the full case studies are available from the authors), we turn to the implications of the paper: what was learned from these particular prototype case studies and what one can in general expect to achieve with prototype case studies.

The institutional settings of the two cases have much in common. Each case study focuses on a long-term, Department of Energy—funded research project at a multiprogram 'national' laboratory. While there are ways in which Los Alamos and Brookhaven diverge—for example Brookhaven has been managed by a university consortium and Los Alamos by industry contractors, they are probably more alike than different. Each is among the largest ten federal laboratories, each depends heavily on Department of Energy financing, each has a broad research portfolio spanning much of the physical sciences, each performs research all along the basic-applied-development spectrum. Most important, each is relatively decentralized with program managers and principle investigators conducting research with a good deal of autonomy.

Table 1
RVM research procedures

Develop sequential, but nonlinear, model(s) of the flow of knowledge from research to exhaustive set of outcomes.
Develop propositions about causal factors related to those outcomes.
Develop indicators of costs and benefits from projects and project-related outcomes.
Select cases on the basis of factors specified in models and hypotheses.
Gather data on cases.
Organize data by writing traditional case studies or by assigning to categories.
Validate data coding conventions (e.g. inter-rater reliability indices).
Use resulting quantitative data in connection with models, determining (through contingency analysis or dynamic programming) the relation of independent variables to information flows, project outcomes, and benefits and costs.

4.2. Prototype case 1: superconducting wire and magnet technology at Brookhaven National Laboratory

Reviewing the lessons from the prototype cases does not require a complete recapitulation of the case studies. In the interest of space, only two cases are examined here and they are provided in summary form. As part of the RVM project, a prototype case was undertaken in 1995–96 examining a superconducting wire project at Brookhaven National Laboratory (Bozeman & Donez, 1996). A second prototype case below, for a different project was conducted in 1997 examining thermoacoustic heat engines at Los Alamos National Laboratory (Klein, 1998).

4.2.1. BNL project focus[3]

In the early 1970s, researchers at Brookhaven National Laboratory (BNL) began their quest for technically-feasible, economically-viable superconducting wire. In terms of the knowledge developed from the research program, the results are an unqualified success. The research resulted in a great many publications and even scientific awards. However, the research did not lead, at least not directly, to a commercially successful superconducting wire and, even now, there is no sign that the research will contribute to product development or to other economic spill-overs. But the case study is a story of *indirect* and evolutionary impacts of basic research. Despite the fact that the interesting scientific results generated in the superconducting wire case never yielded significant economic or social benefit, the line of research enabled progress in an ancillary applied science project- research and development of advanced magnet technology.

Among the many substantive (as opposed to research design or methodological) lessons gleaned from this case, the following stand out:

- The harnessing of basic and applied research in the same or proximate projects has salutary effects for each.
- Basic research has the potential to contribute *short term* to a variety of technology development paths, including ones not easily anticipated.

[3] The data for this case is derived from interviews conducted by Barry Bozeman and Francisco Donez at Brookhaven National Laboratory on 27 March, 1995 and subsequent telephone interviews. Bozeman is the case author. Brookhaven interviewees included Dr James Davenport, Chair, Department of Applied Science, Dr Mas Suenaga and Dr David Welch. Others interviewed included Dr Andrew Kin, Technical Director, Magnetics Division, Crucible Research; C. David Claspell and Dr V. Prem Panchanathan, Senior Engineer, Delphi Energy and Engine Management Systems, formerly Magnaquench.

- The instrumentation and equipment needs for basic research can themselves provide economic benefits.

4.2.2. Case summary

The U.S. Department of Energy's (1994, p. 12) summary of materials sciences programs lists two related projects at BNL, each under the leadership of Dr Masaki Suenaga. Project No. 53, 'Superconducting Materials',[4] examines the properties of high temperature and critical field superconductors and Project No. 54, 'Basic Materials Science of High-T_c Conductor Fabrication',[5] focuses on problems associated with the fabrication of conductors for magnets and power using High-T_c superconductors. But summarizing the research in terms of the boundaries specified in one fiscal year's report does not do justice to the BNL stream of research on superconducting wire, materials and magnets. It is not easy to even conceptualize 'the project' and, indeed, when we tried to describe our project focus to Dr David Welch, one quite familiar with the BNL work stream in this field, he had difficulty determining when one 'project' began and another started. Clearly, to the scientists working in this area at BNL the research is not viewed in distinct project packages but as a constantly evolving, inter-related set of research problems and approaches.

The projects listed in the 1993 funding portfolio evolved from work begun at BNL in the late 1970s. The earlier work was focused on developing superconducting wires. The technical problem was the recognition of the need for less brittle filaments. The superconducting wires were not useful unless very fine and the finest of wires, chiefly niobium or titanium, were also quite brittle.

In the early 1970s, BNL work was on niobium and tin and involved manufacturing of rods, pressing and extruding. A major target application was commercial nuclear reactors. During the 1970s some funds were available from one of the major contemporary research programs at BNL during the 1970s and 1980s, research on superconducting transmission lines.

The momentum for the superconducting wire phase of the research began to wind down about 1985–86 and the researchers began to look for other related projects that seemed to hold promise. But during the 15 years or so of the superconducting wire project, a wide variety of scientific papers was produced and there was considerable progress in basic research on the physical properties of superconductors. The chief commercial impact of the work was providing U.S. business with the capability to make niobium—tin superconductors less expensive and having desirable physical properties compared to niobium—titanium, the previously used material.

One of the reasons for the technical success of the superconducting wire projects, according to Dr Suenaga, was the mix of research specialists working in the same building. During most of the project there were three groups, an accelerator group working on dipole magnets for future accelerators and using ductile superconductors; a solid state physics group doing basic research on superconductivity and; the primary materials group (Dr Suenaga's group) applying electron microscopy to materials characterization. The BNL materials are not useful when there is a need for wires of great length. Today, the more common approach is to use central core of niobium/tin wire as a solid, heat treat it, and form a compound from the result. The process and materials pioneered at BNL are not currently in widespread use.

In sum, the BNL work on superconducting wire yielded considerable scientific knowledge and entailed some practical application of the materials and processing technology developed, but seems to have been, from one perspective, an interesting 'path not taken'. As so many energy-related technologies, superconducting wire seemed to have less potential after the oil shocks and also when there was less reluctance to place transmission wires overhead, but changing times and changing needs lead to new opportunities at the same time as old ones seem less promising.

4.2.3. The rest of the story: magnet technology

By the mid-1980s the BNL research program on superconducting wire had dissipated and, indeed, the level of general activity on superconductivity was beginning to diminish. However, a congeries of events, including declining DOE support for superconductivity, the 'end of the line' on both the superconducting wire and superconducting transmission line projects at BNL and an explicit interest in public policies supporting federal laboratory–industry interaction, came together in the late 1980s, resulting in a research program focused on developing commercially viable magnet technology.

Dr David Welch described the transition as a sort of 'internal technology transfer'—the skills developed from long experience with one line of research were transferred to another. The permanent magnet work began slowly in the mid-1980s with Dr Welch, modest internal laboratory funding and one post doctoral researcher. After obtaining significant results, the funding was expanded for new initiatives.

General Motors was a particularly prominent industrial collaborator on the permanent magnets project.

[4] Investigators include M. Suenaga, B. Budhani, D. Welch and Y. Zhu. The research was funded during the fiscal year at the level of $1,139,000. During the fiscal year reported here, the research focused, specifically, on "...theoretical models of interatomic forces, lattice defects, and diffusion kinetics in superconducting oxides, ...defects in superconducting compounds, flux pinning, properties of composite superconductors" (DOE, 1994, p. 12).

[5] Investigators include M. Suenaga, Q. Li, and Y. Zhu. The research was funded at the level of $608,000. The specific focus was on "characterization of microstructural and electromagnetic grain boundaries (in ceramic materials)...in order to gain increased understanding of the nature of coupling" (DOE, 1994, p. 12).

General Motors sold its magnet producing division several years ago because of problems in developing skills and capacity to work with new materials required for the latest generation of magnets. Ten years ago, nidimium iron boride was first discovered. The auto industry is a major consumer of magnets for motors and actuators and the new material was appealing. GM decided to make its own permanent magnets for use in manufacture.

The interaction between GM (in its Magnequench division) and BNL began in early 1993. Magnequench had already begun, in 1992, working with Idaho National Engineering Laboratory and, thus, had some experience working with federal laboratories. General Motors supplied BNL with materials because the researchers at BNL could provide some expertise and equipment not available at GM. Dr Mas Suenaga and Dr Welch then sent a research proposal to the Department of Energy's Basic Energy Sciences Division to work collaboratively with Magnequench. The proposal was funded. The work, begun in 1993, is still very much in progress.

More powerful magnets are attractive to the auto industry because the more powerful the magnets the smaller and more efficient (energy per unit volume) the motor. A few years after the discovery of the new magnet material, GM built a plant and division to produce the magnets. The Division, Magnequench, had its plant built in Anderson, Indiana. The first Magnequench product, a neodymium magnet, was introduced into the market in 1985.

About three years ago, GM decided to sell Magnequench and now is in the process of doing so. Magnequench is being sold to MQI, a U.S.-based conglomerate funded by a Chinese group and New York City investment firm. The decision to sell Magnequench was part of an overall strategy of GM to focus on its core businesses.

4.2.4. Impacts

Thus far, the chief impact of the collaboration has been to help Magnequench obtain more applied scientific knowledge relevant to the development of new products. The research aimed at product development focuses chiefly on materials characterization and ways to optimize for magnets. The class of magnet at which they are aiming is not the highest performing, but one which will have improved magnetic qualities at a reduced price. The magnet would be used in appliances and autos but not for such high technology applications as computers. The next stage of work involves analysis of the optimal composition of the alloy, testing and looking for modifications.

Production is not just around the corner. The introduction of prototype products based on this research is unlikely to occur before about 1998. Then another two years or so would be required before the actual start of production. Another consideration is that about four years are needed to get on design cycles of autos and about 1.5 for appliances.

The stakes are considerable. The market for magnets with the type applications pursued in this project is about \$600 million. If this product is successful it might have about 10% of that market.

Table 3 provides a summary of the impacts provided thus far in this collaborative project. The chief impact is in enhancing the company's (and BNL's) knowledge and technical capability. The project has also influenced Magnequench's research agenda and has provided some of the information needed in the process of developing a commercial product.

4.2.5. Policy conclusions for BNL cases

These linked cases, the superconducting wire project and the permanent magnet project have provided considerable insight into the process by which fundamental knowledge contributes to practical outcomes. *Without the basic research performed on superconductivity, BNL would not likely have had the knowledge and equipment to contribute to the magnet development project.* As a result of years of progress in basic research, the ability of BNL to contribute to practical problem-solving has been enhanced.

The ability of BNL to work on practical, commercially-relevant projects was, to a large extent, an unplanned outcome of BNL's patterns of organization and collaboration. As several people at BNL noted 'we're in the same building'. The fact that theoretical physicists, experimentalists, materials scientists and electron microscopists interact routinely vastly improves the chances that BNL personnel will be able to deploy pure research in the service of applied problems and, ultimately, commerce.

Industry is attracted to BNL because of the unique skills of BNL scientists but also because of available equipment. None but the largest companies has sufficient capital to invest in the equipment required to work at the frontiers of materials science. Not even the largest companies invest in such enormously expensive equipment as synchrotrons.

The case illustrates the extremely long development time required after the federal laboratory has made its contribution. The amount of testing and development work still needed for developing new magnets requires several years, but there is the added requirement of fitting into product life cycles of the respective products.

Finally, the case shows that it is *not always easy to allocate 'credit' among the various federal laboratory contributors*. In the case of Magnequench, the original working relationship was with INEL and they have continued to work with both INEL and BNL. At the same time, Crucible Metals is also forging ties to Oak Ridge National Laboratory and has already begun working with

Argonne and INEL. This raises a policy question as well as a case research question—is there a need for a scientific and technical 'traffic cop' to route federal laboratory contributions, or is it better to rely on the initiative of the companies and the federal laboratories even if some inefficiencies result? Perhaps allocation of credit among the particular laboratories is less useful than determining the contribution of the federal laboratory apparatus.

4.3. Prototype case 2: thermoacoustic engines at Los Alamos National Laboratory[6]

The technology of heat engines underlies many practical devices, including automobile motors and air conditioners. Although efficient in their use of energy, today's heat engines' use of moving parts renders them expensive and maintenance-intensive. Research at Los Alamos National Laboratory that is funded by the Department of Energy's Office of Basic Engineering Sciences (BES) could change all that.

Thermoacoustic heat engines and refrigerators currently being developed at Los Alamos National Laboratories employ no moving parts. Unlike the conventional designs the thermoacoustic devices harness the temperature, pressure, and displacement oscillations in a sound wave to achieve heat flows. If successfully developed a thermoacoustic heat engine would be cheap and easy to manufacture and would offer a high degree of reliability in operation.

The history of BES-funded thermal physics research and its applications falls into two parts. The first is the pioneering work in low-temperature physics done by John Wheatley. The second is the subsequent research and development in thermoacoustics that grew out of his work.

Although only 59 years old when he died unexpectedly in 1986, John Wheatley was recognized as one of the great low-temperature experimental physicists of the 20th century. His study of the superfluid phases of the rare He-3 isotope of helium was considered one of the most significant topics in low-temperature physics in the 1980s, and colleagues considered him to have narrowly missed a Nobel Prize award. Wheatley also had a strong interest in the application of his research. He founded a company to bring some of his ideas into applications and he shifted from university to national laboratory and back again as his focus shifted between research and technology development.

From 1966 to 1981 Wheatley worked at the University of California (UC) at La Jolla where he worked on superfluid helium-3. In his final years there, he became interested in thermoacoustic engines. These would be a new development and would be sufficiently rich and complex to challenge one's understanding. Furthermore, the promise that something practical could result from an engine with no moving parts appealed to him.

In 1981 Wheatley moved to Los Alamos National Laboratory (LANL). There his research was divided into four areas: heat engines that used liquid instead of gas (to avoid condensation problems), low temperature hydrogen, thermoacoustic engines, and thermal convection.

In 1985 he returned to the University of California, this time to the Los Angeles campus. With a joint appointment with Los Alamos, he retained some ties to LANL. Shortly after his return, however, he died.

After Wheatley's departure and death there was uncertainty at LANL whether his work would be carried on and, if so, how. There were a number of students, post docs, and staff working on projects whose status became uncertain. However, BES gave the researchers continued funding, and laboratory scientists picked up advising duties of students and post docs. Work continued. Wheatley's work in thermoacoustic engines was taken over by Greg Swift, a physicist who had come to LANL as a post doc in 1980.

Swift continued the thermoacoustic research in its applied direction, seeking funding for his work from a variety of sources. An early funding opportunity came from the Strategic Defense Initiative, which provided money to develop a cooling system for infrared sensors on satellites. Although successful in its outcomes, this project fell victim to a change in priorities in SDI away from longer-term research and toward near-term applications. With the end of the cold war, defense funding grew less abundant.

As concerns with international competitiveness grew, funding for commercial research increased. A new opportunity to continue refrigerator development soon arose. In 1992 the Tektronix Corporation expressed an interest in developing the LANL technology for application as a compact and reliable cooler for cryogenic electronics. Tektronix and LANL joined in a cooperative research and development agreement (CRADA) and received funding from the DOE's Technology Transfer Initiative (TTI) as well as from Tektronix. During the next two years they made further improvements in the technology to reduce its size and raise its efficiency. However, in 1995 Tektronix changed its strategy and de-funded the project.

Around this time another funding source became available. Cryenco, a small firm in the liquefied natural gas business, began funding development in 1994 in the hopes of creating a low-cost, high-reliability liquefier of natural gas. Working closely with Cryenco, Swift's group began developing a prototype natural gas liquefier with a capacity of 500 gallons per day. If successful this tech-

[6] The data for this case is derived from interviews conducted by Hans Klein and Barry Bozeman at Los Alamos National Laboratory on 17 August 1997 and subsequent telephone interviews. Klein is the case author. Los Alamos interviewees included Greg Swift and Bob Ecke.

nology could have a major impact on the natural gas industry, allowing much smaller liquefication facilities than are economically feasible today while reducing construction and operation costs.

Just as Wheatley had gone to LANL to pursue his research in an applied direction, so Swift went to Cryenco to carry development to industrial application. Swift worked on site at Cryenco's Denver headquarters for nearly twelve months in 1996 and 1997 building a prototype device. Although continuing his affiliation with LANL, he had to physically pursue development in the context of the industrial application site.

In summary, Wheatley and Swift carried development forward through a variety of institutions and a variety of projects. Research was conducted at University of California, then Los Alamos National Laboratory, and then at the Cryenco Corporation. The thermoacoustic heat engine was reinvented as a defense technology, a commercial technology for instrumentation, and an energy technology.

In its dynamic institutional and funding environment, the funding received from DOE/BES provided valuable stability. BES funding served as a 'flywheel' to maintain research activities even when external resources were scarce. As R&D went from basic science, to development, to industrial application, BES funded John Wheatley and Greg Swift in a variety of activities and in a variety of institutional settings.

It would appear that the personal drive and entrepreneurship of both Wheatley and Swift figured prominently in the program's success to date. That success attests to their desire and ability to overcome institutional barriers and categories in the research establishment. When Wheatley took his research in an applied direction he came to Los Alamos to find a supporting setting. When teaching beckoned again, he left and returned to the university. No institutional setting was adequate for all that he wished to do, but he succeeded in moving between institutions in pursuit of his goals. Similarly, Swift's entrepreneurship kept development moving forward despite a series of setbacks. He successfully hopped from one funding source to another, always moving technology development forward.

BES funding played a vital role in supporting this entrepreneurship. In a context of uncertainty, it provided stability and security. When resources ran out, it served as flywheel to carry research through the dry periods.

5. Lessons learned from the prototype case: lessons for the RVM project

The researchers' objective for the prototype cases was to inform their own planned research and, thus, many of the 'lessons learned' given below are directly relevant to the RVM project. Some have relevance to the general issue of learning from prototype cases. A conclusions section address more generally the issue of learning from prototype cases.

5.1. The project as a unit of analysis

In a previous application of RVM to technology development projects of the New York State Energy Research and Development Authority, the project was taken as the unit of analysis and the project-focus presented little difficulty. In the BNL prototype case, it seems clear that projects evolve from earlier projects, often in a way that makes the project unit of analysis highly artificial, and are intertwined with projects funded by others. The problem is that the boundaries of the research activity do not coincide with the administrative boundaries required by the organizations providing support. This problem can be ameliorated by not being restrictive about the unit of analysis, focusing on programs and multiple projects when appropriate, single projects in other instances. Such an approach defies the usual conventions of sampling design, but in the application of RVM the choice of cases is based more on theoretical supposition and case selection models than on conventional sampling theory.

5.2. Multiple actors at multiple sites

In some instances industry partners work simultaneously with more than one laboratory. Most BES-funded projects are performed within a single laboratory (though there may be others with BES funding working on similar problems at other laboratories). This implies that sampling must not be restricted to single projects from single laboratories. Cases should adequately reflect the participants in the research.

5.3. Projects spanning multiple sites

In interviews the scientists at LANL emphasized the institutional factors affecting their work. Different institutions had different strengths and weaknesses, and so scientists had to change their institutional settings in order to pursue different tasks in the overall R&D process. For example, in the institutional setting of a university, scientists could work on basic research but were discouraged from developing their findings into applied technology. In the institutional setting of a national laboratory, scientists could engage in technology development but received less support for basic work in theory development. Finally, in order to advance an early technology to application, scientists had to work in residence at an industrial site for long periods of time. The different tasks that characterized research, development, and application could only take place in different institutional settings.

This suggested that future case studies would need to investigate more than one institutional setting per project. Multiple sites would have to be visited for a single project. Moreover, particular attention would have to be paid to the crossing of boundaries between sites. Thus an initial research strategy of performing one site visit per project was replaced with a strategy of multiple site visits per project.

5.4. *BES funding and contributions of others*

The sponsor for the RVM project is the Department of Energy's Office of Basic Energy Sciences and, thus, the case focus is on the RVM sponsor's projects. But in many cases, BES funding is a poor criterion for analytical focus. Often BES funding co-exists with other sources, supplements pre-existing funds, or diminishes when the more applied aspects of projects come to the center. A partial remedy to this problem is to take BES funding as a beginning point but identify other sources of support, using composition of BES funding as an 'independent variable'. In other words, the proportion, timing, and phase of BES funding can itself be used as a predictor of impact.

5.5. *Qualitative differences in funding*

As revealed by the LANL case, successful research often reflected successful strategy. Good scientists could 'read' their institutional context and move around in it as their needs changed. Thus we find a physicist moving from the University of California to Los Alamos National Laboratory in order to pursue a new research direction in refrigeration. He could 'read' university institutions and recognize that they would not support this applied turn in his work as well as would a national laboratory, so he transferred to a national laboratory. Because he had an institutional strategy and was able to execute it, he could hop from one institutional setting to another.

Recognition of the role of strategy illuminates qualitative differences in research funding. Some research funding was linked to specific institutions, some was linked to specific projects, and some was linked to specific individuals. BES funding was mostly linked to individuals. BES continued to fund a scientist as he moved between different institutions, allowing him to pursue a line of research in a variety of settings. This funding approach of picking an individual scientist and funding him or her for many years gave scientists some independence from their institutional setting and supported them when make inter-institution or inter-project transitions.

As a result of this observation researchers refined their concepts. The simple variable 'funding' was elaborated into different kinds of funding, depending on whether it was linked to individuals, institutions, or projects.

5.6. *Long time-lines*

Probably the most familiar obstacle to assessing basic research outcomes is the extremely long development time required after the basic researchers have made their primary contribution. In the BNL case, the amount of testing and development work still needed for developing new magnets required several years and there was the added requirement of fitting into product life cycles of the respective products. In addition, there is the obvious point that benefits often unfold at a relatively slow pace, especially for projects at the very end of the basic research spectrum. Cases taking the project/program as the sampling unit cannot easily cope with the complexities of basic research time lines. Thus, we suggest a complementary sampling strategy, one to be undertaken in addition to (rather than in lieu of) traditional sampling. In at least some of the cases, the starting point for a case should not be the initiation of a project or program but some technological or commercial outcome that is (hypothesized to be) a the result of the BES-funded project or program. In many instances, the 'end of the story', or at least the last chapters, can be the point of entry into a case. Not all case studies need follow the same chronological entry point.

5.7. *BES-funded contributors and industry contributors*

Sometimes it is virtually impossible to sort out the contributions of BES-funded participants from those of their industrial partners. In some instances the industrial partners and government or university-based researchers are working side-by-side, even producing co-authored research studies. There is no easy solution to this 'problem'. One approach (as above) is to treat the mode and level of participation by collaborators as an independent variable.

5.8. *Proprietary considerations*

In some cases the assessment of commercial benefits is likely to require information which companies hold as proprietary, usually as trade secrets. It is difficult to determine from the three cases provided here just how often proprietary consideration will inhibit the ability to do case studies. But in such instances, cases can be nonetheless valuable; since the object of the cases is not so much to determine the exact monetary benefits, the proprietary information is of less interest than the processes which led to output of proprietary value.

5.9. *The 'path not taken': assessing the value of dead ends*

It is widely observed, and to a large extent correct, that projects that seem to be dead ends often make a contribution by eliminating dead ends. This does not seem to be an obstacle to the conduct of case studies. Indeed, it is certainly useful to consider some dead ends as a point of comparison with high impact cases. The previous application of RVM intentionally included some dead end projects for a point of comparison. However, results indicated that some projects believed to be dead ends were not at all (just as some believed to have high impacts turned out to have negligible impacts). Given the experience in these previous cases as well as from the current study, some projects should be included that are initially believed to be low impact or dead end projects. As a minimum, such projects will yield useful comparative information. Future cases should seek to determine the value of dead ends and whether results of dead end projects are known to others, perhaps preventing them from wasting resources on approaches that are not scientifically, technically or commercially viable.

5.10. Commercial benefit as 'opportunity cost'

Project-level cases are not adequate for more general assessments of the opportunity costs of resources invested in commercial work, but case studies of projects can benefit from the use of 'opportunity cost' assumptions. It is useful to at least maintain the perspective that there are alternative uses for resources directed to commercial work and, moreover, it may be possible in particular cases to determine alternative uses of resources within the focal project. That is, there may be decision foci in which commercial pursuits are clearly traded off against other uses of resources and, just as important, there may be instances where commercial benefits do not exhaust alternative uses of resources (e.g. when the scientific goals for a project are entirely compatible with the commercial goals.

5.11. Termination of benefits and impact assessment: how much to claim?

One attribute of basic research that is especially challenging for one wishing to assess its impacts is determining the 'cut-off point' for benefits. If research is truly fundamental it often has fundamental impacts, ubiquitous in time and domain. At some point, however, it becomes implausible to attribute additional benefits to one research base; even if a research outcome seems a necessary condition for an impact, the outcome may well have required a hundred or a thousand additional necessary conditions. While it is vital to resolve this problem for any monetary accounting of costs and benefits of research, the RVM approach does not require such a level of measurement precision and, thus, the problem is not intractable. In such cases, the model will simply be so complex as to have little practical value and this complexity is itself a boundary setting device.

5.12. Limits of recall: a particular problem of basic research assessment

All case studies must face problems of respondent recall and the various biases (e.g. telescoping, post hoc rationalization) associated with recall. In the case of basic research assessment this is a particular problem because there are competing values. On the one hand, one wishes to choose a project/program that has had sufficient incubation period. On the other hand, one wishes to have interviewees who are still alive and can remember events associated with the research. While there is no good solution to this problem, one rule of thumb is to simply ensure considerable variance within the practicable time band. For example, one may wish to posit the assumption that at least one year is required for results and that projects more than ten years old are not feasible from the standpoint of recall (though in the latter case it may sometimes be possible to substitute documents and existing historical evidence for direct interview evidence).

6. RVM update: using the prototype case studies

By this point (summer 1998), most of the RVM case studies have been conducted and the cases have, as intended, been used as basic material for RVM maps and taxonomies. While not all the lessons learned in the protytpe case studies have been directly implemented in the RVM case studies underway during 1997–98, most have. As a direct result of the earlier case studies the current RVM project has taken the following steps:

(1) The project has been abandoned as the exclusive unit of analysis for the case studies.
(2) The 'one year to ten year' (at least one year, no more than ten) rule has been used in case selection and interviewing.
(3) Cases have been included and interview protocols developed to consider explicitly apparent 'dead ends'.
(4) Opportunity cost hypotheses are developed, examining the impact of commercial work on ability to perform basic science.
(5) Cases have been developed where the Office of Basic Energy Sciences is not the primary source of funds.
(6) The very nature of 'benefit' and 'impact' has been reexamined and a new approach to determining value has been developed (Bozeman & Rogers, 1997).
(7) Investigation has focused on how different kinds of funding affect outcomes.
(8) Greater attention has been given to the flow of knowledge as embodied in individuals as opposed to formal sources of information such as journal articles and patents.

7. Learning from the prototype cases: some general issues

The most important point about the use of prototype cases is that their utility must, first and foremost, be judged against the specific needs of the project which they are designed to inform. Thus, the value of prototype cases will, in part, be a function of the researcher's ability to reflect on the critical dimensions and needs of her project. The foregoing analysis of lessons learned for the RVM project is illustrative of some of the uses to which prototype case may be put. But, in this section, we consider some more general implications of the use of prototype cases.

Regardless of the research methods employed *after* the implementation of the prototype cases, these initial cases can have strong heuristic value. Among others, prototype case studies can shed light on the following questions:

(1) *What is the appropriate boundary for the study?* In aggregate data studies, the boundary question often is easier, or at least seems easier, because one is working with sampling logic. This is possible with case studies but presents many problems when the population is not clear or is so small that sampling logic makes little sense, then boundary-setting becomes especially challenging. While it is sometimes possible, especially from a grounded theory perspective (Glaser & Strauss, 1967) to make design decisions 'on the fly', a prototype case often helps in bounding the phenomena of interest.

(2) *Which unit of analysis?* Choice of unit of analysis often is among the more difficult research problems confronting the case study analyst. The prototype case studies described here were important in sensitizing the researchers to the limitations of the project as unit of analysis. It seems to us that prototype case studies can serve a similar purpose in most cases, at least so long as the researchers come prepared with the right questions. The 'right' questions pertain to the purposes and hypotheses of the particular study. In general, the analysts should always attend to the relation of the information being obtained from sources to the research questions of interest. If there is a mismatch it may be due to a focus on the wrong unit of analysis or, more common, on an insufficient number of units of analysis.

(3) *What is the best 'time slice'?* A special boundary-related problem, one easily recognized in case studies but one affecting most research methods, is the appropriate time focus. What slice of behavior during what time period does the research examine? A benefit of systematic consideration of this question is that it often provides new insights into the slippery nature of causal inference. What looks like a 'cause' at one time, may seem a correlate or a pre-condition at another.

(4) *Where is the data located?* Researchers have to conduct site visits where the data is located. A prototype case can reveal that a single case may require visits to more (or fewer) sites than researchers originally anticipated. A change in the number of site visits can have large impacts on budgets and field work strategies.

(5) *What are the prospects for data access and collection?* Often the researcher has only limited knowledge of the practical difficulties entailed in obtaining data. A prototype case often requires the researcher directly to confront the gap between ideal and the attainable data and, in the process, rethink the objectives and expectations for the research.

(6) *What hypotheses are suggested?* Not every researcher begins with systematically formulated hypotheses and, indeed, some qualitative researchers eschew formal hypotheses. However, in studies, such as our RVM project, in which there is an interest in developing preliminary hypotheses, prototype cases can be extremely valuable. In general, the use of case studies for hypothesis and theory development are well-documented (e.g. Eisenhardt, 1989). A special problem in the use of case studies is developing insight into which apparent causal paths are unique and which generalizable. Prototype cases can provide insights as to the relation of unique to patterned phenomena.

(7) *What about 'instrumention effects' due to case analyst variance?* The prototype case studies presented here began with quite similar purposes, were conducted in similar setting, used similar protocols, included one person (Bozeman) common to each case study. But the two prototype cases have different 'flavors' owing to the interests, backgrounds and focus of the respective authors. Indeed, that is itself a lesson: the difficulty, perhaps the impossibility of using multiple case authors and expecting similar foci and results. No amount of planning and preparation will 'undo' the interests, expertise, and styles of particular case writers—nor should it. Each case writer has something unique to bring to the case. In experiments, we refer to this as an 'instrumentation problem' and seek to standardize experimenters' interactions with subjects. But this 'problem' in case studies is a natural outgrowth of an approach that is as much interpretation as science.

(8) *How does one cope with demands of the particular field research setting?* In field research, each setting provides its own problems, some of which are easy to anticipate, others of which could not possibly be imagined. For example, in the federal laboratory setting, one knows ahead of time that there are security measures in place and one plans for these. But other insights into the setting are difficult to obtain by any other means than a prototype study. Is it best

to interview supervisors and subordinates apart or together? If laboratories crowded by equipment, are there convenient electrical outlets for laptop computers? In interviews are scheduled at lunch, is it possible to sit in a quiet corner of the institutional cafeteria or is there no quiet corner? Often prototype case studies can help with questions of setting, both the mundane (but practically important) issues and the issues vital to method and theory.

(9) *What is it like from the other side of the mirror?* Often, the researcher has only a limited perspective on the range of subjects' views about the research and the researchers. In many instances, prototype cases quickly and easily provide this perspective. Sometimes hostility and suspicion are rampant, other times respondents are comfortable with the research and only interested in being helpful. In more complex cases, subjects have their own agendas for the research, sometimes difficult to ferret out. Prototype cases are almost always useful for understanding the subjects' views about the research.

8. Deploying the prototype case

The questions are examples of some of the uses of the prototype case studies. The potential uses are nearly as broad as the set of researchers' needs. When one has the luxury of using case studies to inform a larger project, the prototypes, if properly deployed, can prove highly beneficial. What are 'properly deployed' prototype cases? First, they should explicitly examine phenomena of interest in the broader research project (obvious, but the closer the better) in a setting closely resembling that for the broader research. Second, pain staking records are required because at the beginning of the prototypes it is not at all clear which 'lessons' will emerge or become the most important. Third, the researcher must be open to thoroughgoing reconsideration of the basic theoretical and methodological assumptions of the project (as, in the RVM case, even the very notion of 'value'). Fourth, when there is only one prototype case, the researcher must constantly attend to the question 'is this finding (problem, implication) unique, or is there reason to believe it may obtain for other cases or data'?

Despite the many advantages of prototype cases, there are instances in which the approach is not suitable. When the research design requires that the exposures of the subjects or the researchers must be 'fresh', prototype cases 'contaminate'. If the phenomena is so sparsely distributed that even a single case cannot be used solely for learning purposes, then the prototype is not desirable. Most important, or at least most common, the prototype case may not be a good investment in terms of its cost to the research project. Case studies are never inexpensive and for projects on a small budget, the prototype may cost as much as the project itself. But in those cases where resources are adequate and when the intervention does no harm to the objectives of the project, the prototype case is likely to prove a good investment, returning hypotheses, improvements in craft, and phenomenological understanding.

References

Bailey, M. (1992). Do physicists use case studies? Thoughts on public administration research. *Public Administration Review*, *51*, 47–54.

Bozeman, B. (1997). Progress report, No. 2: The R&D value mapping project. Report to the Department of Energy, Office of Basic Energy Sciences. Atlanta, GA: School of Public Policy, Georgia Institute of Technology.

Bozeman, B., & Donez, F. (1996). Brookhaven National Laboratory, Superconducting materials and magnet technology. Case study prepared for Sandia National Laboratories.

Bozeman, B., & Melkers, J. (Eds.) 1993. *Evaluating R&D impacts: methods and practice*. Norwell, MA: Kluwer Publishing.

Bozeman, B., & Crow, M. (Forthcoming). *Limited by design: chaos in the U.S. R&D laboratory system*. New York: Columbia University Press.

Bozeman, B., & Roessner, D. (1995). Prototype case studies for R&D value mapping. Atlanta: Technology Policy and Assessment Program, Georgia Institute of Technology. Monograph prepared for Sandia National Laboratories.

Bozeman, B., & Rogers, J. (1997). Knowledge value communities: the proof is in the 'putting'. Paper presented at the Annual Meeting of the Society for the Social Study of Science, Tucson, Arizona, October 1997.

Bozeman, B., & Fellows, M. (1988). Technology transfer at the U.S. national laboratories: a framework for evaluation. *Evaluation and Program Planning*, *11*, 65–75.

Bozeman, B., & Kingsley, G. (1997). R&D value-mapping: a new approach to case study-based evaluation. *Journal of Technology Transfer*, *22*(2), 33–42.

Bullock, R., & Tubbs, M. (1987). The case meta-analysis method for OD. *Research in Organizational Change and Development*, *1*, 171–228.

Campbell, D. (1975). Degrees of freedom and the case study. *Comparative Political Studies*, *8*, 178–193.

Crow, M., & Bozeman, B. (1987). R&D laboratories' environmental contexts. *Research Policy*, *13*, 329–355.

Dyer, W. Jr, & Wilkins, A. (1991). Better stories, not better constructs, to generate better theory: a rejoinder to Eisenhardt. *Academy of Management Review*, *16*, 613–619.

Eisenhardt, K. (1989). Building theories from case study research. *Academy of Management Review*, *14*, 532–550.

Glaser, B., & Strauss, A. (1967). *The discovery of grounded theory: strategies of qualitative research*. London: Wiedenfeld and Nicholson.

Jauch, L., Osborn, R., & Martin, T. (1980). Structured content analysis of cases. *Academy of Management Review*, *17*, 517–525.

Kennedy, M. (1976). Generalizing from single case studies. *Evaluation Quarterly*, *3*, 661–678.

Kingsley, G. (1993). The use of case studies in R&D impact evaluations. In B. Bozeman, & J. Melkers, *Evaluating R&D impacts: methods and practice*. Norwell, MA: Kluwer Publishing.

Kingsley, G., & Bozeman, B. (1997). Charting the routes to commercialization: the absorption and transfer of energy conservation technologies. *International Journal of Global Energy Issues*, *9*, 1–2.

Kingsley, G., & Farmer, M. C. (1997). Using technology absorption as an evaluation criterion: The case of a state R&D program. School of Public Policy Working Paper 97-02. Atlanta, GA: Georgia Institute of Technology.

Klein, H. (1998). Los Alamos National Lab: Thermal physics. Case study prepared for Department of Energy.

Larsson, R. (1993). Case survey methodology: quantitative analysis of patterns across case studies. *Academy of Management Journal*, *36*, 1515–1546.

Lee, A. S. (1989a). A scientific methodology for MIS case studies. *MIS Quarterly*, *13*, 33–52.

Lee, A. S. (1989b). Case studies as natural experiments. *Human Relations*, *42*, 117–137.

Lee, A. S. (1991). Integrating positivist and interpretive approaches to organizational research. *Organization Science*, *2*, 342–365.

Luthans, F., & Davis, T. R. V. (1982). An idiographic approach to organizational behavior research: The use of single case experimental designs and direct measures. *Academy of Management Review*, *7*, 380–391.

McClintock, C., Brannon, D., & Maynard-Moody, S. (1979). Applying the logic of sample surveys to qualitative case studies: the case cluster method. *Administrative Science Quarterly*, *24*, 612–629.

Ragin, C., & Becker, H. (1992). *What is a case?—Exploring the foundations of social inquiry*. Cambridge: Cambridge University Press.

Roessner, J. D. (1996). Whisker-reinforced ceramics at Oak Ridge National Laboratory. Case study prepared for Sandia National Laboratories.

Roessner, J. D., Bozeman, B., Donez, F., & Schofield, T. (1996). Understanding thin film deposition at the Stanford Syncrotron Radiation Laboratory. Case study prepared for Sandia National Laboratories.

Stake, R. (1995). *The art of case study research*. Thousand Oaks, CA: Sage Publications.

U.S. Department of Energy (1994). Materials sciences programs, fiscal year 1993. Washington DC: Office of Energy Research, Office of Basic Energy Sciences, Division of Materials Science.

Wilson, S. (1979). Explorations of the usefulness of case study evaluations. *Evaluation Quarterly*, *3*, 446–459.

Yin, R. K. (1981). The case study as a serious research strategy. *Knowledge: Creation*, *Diffusion*, *Utilization*, *3*, 97–114.

Yin, R. (1982). The case study crisis: some answers. *Administrative Science Quarterly*, *26*, 58–65.

Yin, R. (1982). The case study as a serious research strategy. *Knowledge*, *3*, 97–114.

Yin, R. (1994). *Case study research*, *design and methods*. (2nd ed.). Newbury Park: Sage Publications.