Enhancing the Quality of Case Studies in Health Services Research

Robert K. Yin

Objective. To provide guidance on improving the quality of case studies in health services research.

Data Sources. Secondary data, drawing from previous case study research.

Research Design. Guidance is provided to two audiences: potential case study investigators (eight items) and reviewers of case study proposals (four additional items).

Principal Findings. The guidance demonstrates that many operational steps can be undertaken to improve the quality of case studies. These steps have been a hallmark of high-quality case studies in related fields but have not necessarily been practiced in health services research.

Conclusions. Given higher-quality case studies, the case study method can become a valuable tool for health services research.

Key Words. Unit of assignment (the "case"), rival explanations as a design strategy, triangulation of evidence, case study database, case study protocol

Case study methods are being rediscovered in health services research. Much of the contemporary need for case studies is driven by developments in managed care systems that link their multiple components in new ways, producing "mega-systems" of great complexity. Further, the systems' rules are in a high-flux state, continually and rapidly changing. Finally, important corporate affiliations and motivations are extremely difficult to track, much less understand.

All of these conditions favor the use of case studies, over other empirical methods, to gain insight into these mega-systems and to assess them. Rediscovering the relevance of the method overshadows the fact that over the years—and predating the emergence of health services research as a formal professional endeavor—case studies already had been used as a tool to study health services. Examples of such case studies include Leo Srole and his team's descriptive analysis of mental health in the metropolis (Srole, Langner, Michael, et al. 1962); Dorothy Nelkin's penetrating study of a

single methadone maintenance clinic (1973); Richard Neustadt and Harvey Fineberg's study of the national immunization campaign against swine flu (1983); and numerous studies of single facilities such as health centers, hospitals, and community mental health centers (e.g., Bellin, Geiger, and Gibson 1969). In this earlier era, one meta-analysis reviewed the findings from 48 case studies devoted solely to the topic of citizen participation and other modes of decentralized health services (Yin and Yates 1975); another policy series involved 17 case studies of promising medical technologies of the time (U.S. Office of Technology Assessment 1979–1981). Other common topics of case studies in health services research included knowledge and technology utilization; medical and organizational innovations; and the implementation of specific health legislation, policies, and programs, such as Eric Redman's well-regarded case study on the creation of the National Health Service Corps (1973).

From the perspective of case study methods, these examples ably demonstrate desirable features such as the systematic and intense use of archival data (Srole, Langner, Michael, et al. 1962; Bellin, Geiger, and Gibson 1969); insightful and detailed fieldwork by single investigators (Nelkin 1973; Redman 1973); a thorough sifting of policy documents and interview data related to decision making in a nationally prominent setting (Neustadt and Fineberg 1983); and useful frameworks for deriving lessons from multiple case studies (Yin and Yates 1975; U.S. Office of Technology Assessment 1979–1981). From the perspective of health services research, assessments of case study contributions have yet to receive adequate attention. However, the cited examples do suggest the potentially broad applicability of the method—an applicability not limited to the contemporary need for research on high-flux or mega-systems.

Regardless of the perspective, a clear frustration is that the renewed interest in using the case study method in health services research appears not to have acknowledged, much less built cumulatively on, the past. In part to offset the pattern, this article highlights eight desired characteristics of case studies that might be emulated by investigators wishing to conduct new case studies in health services research. Later, four additional features are offered for reviewers of case study proposals, who might be trying to anticipate the quality of a proposed case study before it has begun.

Address correspondence to Robert K. Yin, Ph.D., President, COSMOS Corporation, 3 Bethesda Metro Center, Suite 950, Bethesda, MD 20814. This article, submitted to *Health Services Research* on January 28, 1999, was revised and accepted for publication on June 24, 1999.

DESIRED CHARACTERISTICS OF CASE STUDIES IN HEALTH SERVICES RESEARCH

Using a Design-Oriented Definition of Case Studies

The all-encompassing feature of a case study is its intense focus on a single phenomenon within its real-life context. The method is not troubled by the fact that the context contains innumerable variables–therefore leading to the following technical definition of case studies:

> [Case studies are] research situations where the number of variables of interest far outstrips the number of datapoints (Yin 1994: 13).

The result is a definition that considers case studies from the standpoint of study design and does not automatically associate them with any data collection method (Yin 1994).[1]

In contrast to the role of the context in case studies, experimental designs have the opposite goal: to "control" the context. Survey designs can cover only limited aspects of the context, but they attempt to represent the universe of the phenomenon being studied. Histories do cover both phenomenon and context, like case studies, but are usually studies of the "dead" past.[2]

Furthermore, a related feature is that case studies tolerate the condition whereby the boundary between a phenomenon and its context is not clear. The case study method has sufficient flexibility to cope with this uncertainty, which others have considered an important feature of qualitative research more generally (e.g., Sofaer in this issue). A final real-life complication appreciated in conducting case studies is that both the "case" and its context may be changing over time, adding immeasurably to the number of variables and the complexity of any analysis.

Figure 1 illustrates this definition of case studies by showing three concrete examples as well as by adding yet another feature: that the holistic case (the "hospital," "community," and "managed care system" in the three examples) may contain lesser units of analysis within it. For example, the case study of a single hospital may contain surveys of the clinics and services, staff, and patients within the hospital; however, certain key aspects of the organization as a whole (e.g., the hospital's financial system or its bureaucratic structure) cannot be decomposed.

The illustrations thus also suggest how quantitative as well as qualitative data can be part of the same case study, leading to the likelihood that most case studies will rely on multiple types of data collection. In other words, the technical definition, using case studies as a feature of the design of a study,

Figure 1: Case Study Definition, Illustrated by Three Hypothetical Case Studies in Health Services Research

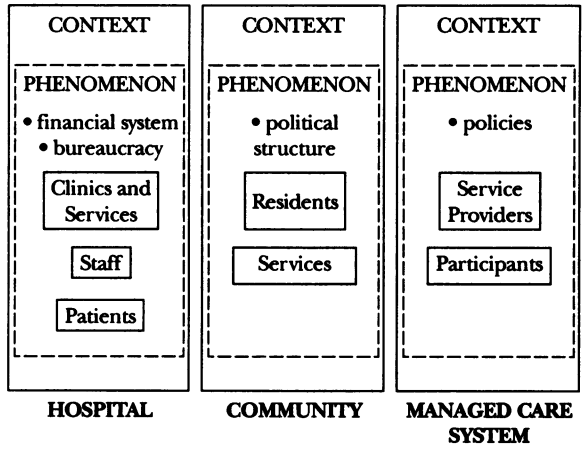

assumes that qualitative techniques and data are part of the case study arsenal. But then, so are quantitative techniques and data. A contrasting framework depicts case study research as a method falling within the arsenal of qualitative research (e.g., Mays and Pope 1996; Clark 1997). This article is directed only to the first of the two frameworks. Investigators working within the second framework will need to seek guidance elsewhere for designing high-quality qualitative research.

Making Generalizations from Case Studies: Replications Based on Theory, not Principles of Sampling

Difficulty in generalizing from case studies has been considered a major shortcoming of the method, whether the research involves single case studies or multiple-case studies. However, the focus on design as the driving definition of case studies provides useful advice for dealing with this problem. The remedy is to consider a case study, as a unit, to be equivalent to an experiment, as a unit; multiple-case studies may then be considered equivalent to multiple experiments. Under this assumption, the problem of generalizing from case studies is no different from the problem of generalizing from experiments—where hypotheses and theory are the vehicles for generalization. To this extent investigators doing case studies are not "theory driven" (a criticism that has been raised by some), but are "driven to theory."

An inappropriate strategy, leading to trouble but nevertheless overcome only with great difficulty, is to use "sampling" criteria in selecting and aggregating individual case studies. Unless a case study is extremely well endowed and is permitted an extensive period of time for completion, the number of cases will always be insufficient to satisfy the sampling strategy. In contrast, by considering each case study to be a separate sub-inquiry, the generalization from case studies can be accomplished by using a replication logic (Yin 1994: 30–32). The needed replication logic can be derived only from hypotheses or theories about the cases. In other words, a theory about what is being studied—and about whether a single case is a "critical" exemplar of that theory or about why some multiple cases might be expected to be replications and others might not—is essential to case study design and analysis.

Figure 2 illustrates both types of generalizing. For "Type One" generalizations, both surveys and experiments may have sampling considerations in selecting respondents or "subjects," but the preferred way to select cases reflects theoretical considerations (e.g., a "critical" case). In "Type Two" generalizations, all three methods—surveys, experiments, and case studies—eventually must cope with policy or theoretical considerations in order to arrive at useful contributions to knowledge.[3]

Figure 2: Making Generalizations: Two Types

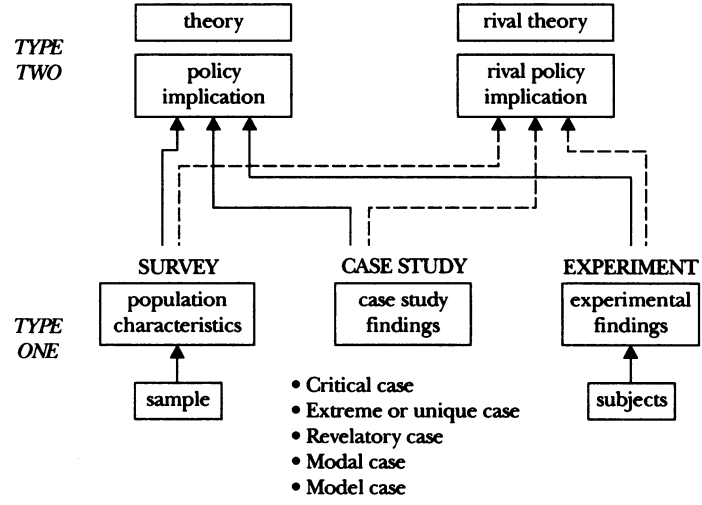

Source: Yin (1994): 31.

Identifying the Unit of Assignment: What Is the "Case?"

Studying a phenomenon in its real-life context still leaves the investigator with the major problem of operationally defining the phenomenon or the "case" being studied. In fact, in doing a case study (or multiple-case studies as part of the same inquiry), the most difficult step is to define the "case" (e.g., Ragin and Becker 1992). Inexperienced investigators are not likely to be sensitive to this difficulty, but it is especially likely to arise in health services research because of the desire to study cases that are rather abstract (one grammatical clue is the use of an abstract, not concrete, noun to define the "case"): "marketplace," "service system," "site," "plan," "comprehensive system of care," and so on. Trying to study any one of these phenomena requires defining them operationally at the outset of the study, even on a tentative basis. Inadequate definition can lead to two problems: (1) the findings might not be about the presumed case but about some other situation; and (2) if multiple cases have been conducted, they might not be comparable in some fundamental way.

The preliminary theoretical propositions used in formulating the case study design provide important guidance for defining the case. Nevertheless, investigators should be prepared for unexpected complications that also may lead to changes in the initial set of propositions. For instance, is the definition of the "case" a health care system at t_1 or after it might have been reorganized at some later time? Similarly, even so seemingly straightforward a case as a "group practice" or a "hospital" may have changed dramatically over time.

As another complication, common language can mislead investigators into thinking that they have satisfactorily defined their case. The words "community" and "site" are such examples, as in doing case studies and conducting site visits to study health care in prespecified communities (e.g., the Robert Wood Johnson Foundation's Snapshots study described in Ginsburg 1996). The Snapshots study did not clearly specify whether the "community" was to be defined as a marketplace, a geographic area, or a set of eligible clients. Site visits were made to the "community" to study changes in service delivery. Complementary surveys were conducted of individuals in the "community" to ascertain whether the changes had affected clients. Little relationship between the site visit and survey findings was observed, and the authors concluded that customers had not yet been affected by the changes in service delivery (see Ginsburg 1996: abstract). However, an alternative possibility is that the individuals who were surveyed were not necessarily part of the service delivery system that was studied in the site visits. The findings under these conditions therefore contain some ambiguity.

Operationalizing the Case Study (e.g., A Logic Model)

Good case studies should contain some operational framework, even if the case studies fall into the classic "exploratory" mode. Even when exploring, some framework should be in place to define the priorities to be explored (e.g., Yin 1994: 28–29).

For most case studies, a common operational framework increasingly takes the form of a "logic model" (Wholey 1979) or a specification of hypothesized cause-effect-cause-effect-cause-effect patterns over time. Typically, a logic model visually portrays these patterns as a path of arrows from left to right and top to bottom, as in Figure 3. The figure represents hypothesized relationships among community networking, community health planning, the production of coordinated health programs, and thus a more effective reduction in the incidence of disease. Having such an operational framework ahead of time helps to define what is to be studied as well as the topics or questions that might have to be covered.

Figure 3: Specifying the Topic Operationally

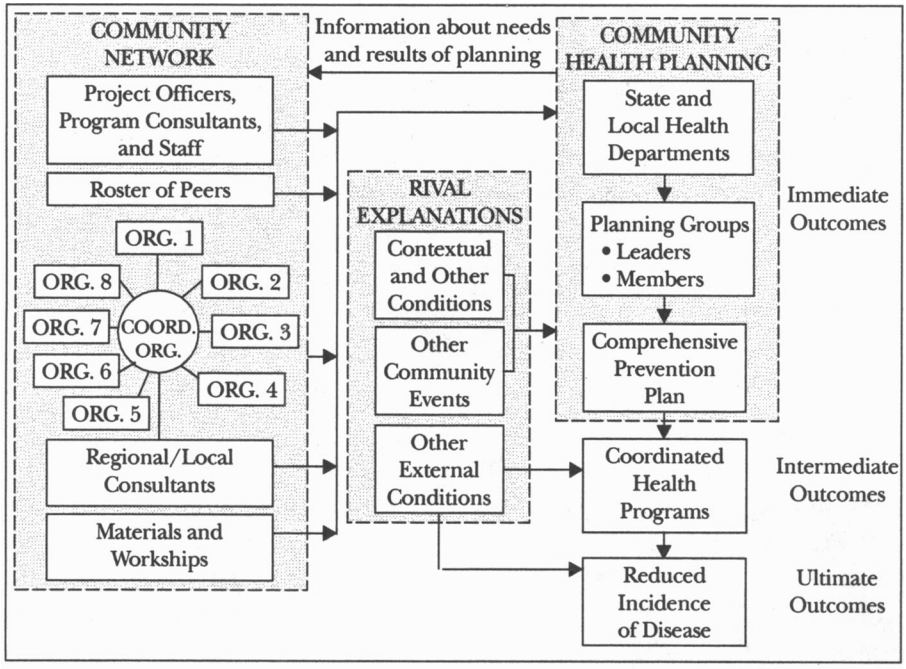

Source: COSMOS Corporation.

Maintaining Discovery and Flexibility

Specifying theoretical propositions and defining operational frameworks ahead of time does not mean that case studies must follow a rigid research design. On the contrary, an invaluable feature of the case study method is the ability to "discover" while in the process of doing the research.

At the same time, many earlier case studies have misused the desired flexibility, leading to accusations of bias and selectivity, which need to be avoided at all costs. The common pitfall is that an investigator will have collected some initial evidence, leading to a desired change in research questions or hypotheses. However, the investigator will then use the same initial evidence as part of the evidence later arrayed to address the revised research questions. Such circularity creates the common criticism of case studies: that an investigator has (unfairly) sought only evidence that supports the initial hypotheses and has not followed a fair procedure in reaching the conclusion.

The recommended way to maintain flexibility and permit the possibility of discovery—while avoiding the trap of conducting biased research—is conceptually simple but not easily (or cheaply) implemented. Nevertheless, the procedure must be followed to accomplish the desired result. For instance, assume that discovery has occurred in a multiple-case study during the conduct of the initial case studies and that such discovery thereby calls into question the initial theoretical framework. The desired procedure is to "recycle" the entire study by revisiting all of the earlier steps (i.e., theory development, case selection, and the design of data collection procedures). These earlier steps need to be revised to accommodate the discovery, possibly leading the investigator to discard some (if not much) of the earlier work. The cycle then needs to be restarted at some appropriate place, now following the ideas emanating from the discovery. Further, the case study that was involved in the discovery should be considered a "pilot" case, and its data should not be used as part of the final array of evidence in the research.

Using Rival Explanations as a "Design" Strategy

The nature of case study research raises yet another frequent dilemma for investigators: how to use experimental or quasi-experimental designs when the number of cases is likely to be small, undermining the ability to form sufficiently large "groups" of cases to suit many such designs. The current state-of-the-art commonly ignores this problem, and a frequent design that is encountered might have three (or four or five)[4] "target" cases and three (or four or five) "comparison" cases.

No alternatives currently exist, but an emerging approach is worth considering. Such an approach puts great emphasis on defining and testing rival explanations as part of the design (Yin forthcoming). The more that plausible rivals are investigated and then rejected, the greater the support for the original hypotheses. Although the level of certainty in the findings will not attain the certainty of true experiments, the main trade-off is that use of the rival explanations method permits the study of social or policy topics that cannot be addressed by experimental designs (Yin forthcoming).

As a result, investigators may want to put serious effort into identifying and collecting data to support or reject plausible rivals, that is, rival explanations from either a logical or empirical standpoint. Key to the procedure is to overcompensate in collecting evidence in support of any rivals—as if deliberately seeking to support them (Patton 1990: 462). With this orientation, investigators are more likely to have their procedures considered fair and unbiased.

Collecting Evidence from Multiple Sources: Triangulation

For case studies, data collection may—and should—involve a broad variety of techniques, not just a single technique such as conducting a site visit. Surveys can be invaluable in gathering case study evidence, but so can archival analyses, documentary searches, and direct field observations. In fact, the more all of these techniques are used in the same study, the stronger the case study evidence will be.

Figure 4 enumerates the main types of evidence and their strengths and weaknesses. A case study investigator needs to be familiar with every technique and should not favor one over the other. The need for such diversity of skills has frequently been overlooked by less experienced investigators, who may incorrectly assume that undertaking case studies is mainly a matter of "living in the field," "telling it like it is," and avoiding numeric data. In fact, a well-executed regression analysis might even be part of a case study. For instance, a case study of a hospital might include quantitative analyses about patients' payments to the hospital as well as a qualitative analysis of the hospital's financial and record-keeping system, to draw conclusions about the effectiveness of the system. In this example, the hospital would be the "case," and conclusions drawn about it would still reflect a single datapoint; the patients would represent a lesser or embedded unit of analysis.

In using multiple sources of evidence, the goal during the data collection process is to amass converging evidence and to triangulate over a given fact.

Figure 4: Six Sources of Evidence: Strengths and Weaknesses

Source of Evidence	Strengths	Weaknesses
Documentation	• stable–can be reviewed repeatedly • unobtrusive–not created as a result of the case study • exact–contains exact names, references, and details of an event • broad coverage–long span of time, many events, and many settings	• retrievability–can be low • biased selectivity, if collection is incomplete • reporting bias–reflects (unknown) bias of author • access–may be deliberately blocked
Archival Records	• [same as above for documentation] • precise and quantitative	• [same as above for documentation] • accessibility due to privacy concerns
Interviews	• targeted–focuses directly on case study topic • insightful–provides perceived causal inferences	• bias due to poorly constructed questions • response bias • inaccuracies due to poor recall • reflexivity–interviewee gives what interviewer wants to hear
Direct Observations	• reality–covers events in real time • contextual–covers context of event	• time consuming • selectivity–unless broad coverage • reflexivity–event may proceed differently because it is being observed
Participant-Observation	• [same as above for direct observations] • insightful into interpersonal behavior and motives	• [same as above for direct observations] • bias due to investigator's manipulation of events
Physical Artifacts	• insightful into cultural features • insightful into technical operations	• selectivity • availability

Source: Yin (1994): 80.

The rationale for such triangulation and the strategies for achieving it are well covered by Michael Quinn Patton (this issue) and therefore are not discussed further here. To be emphasized, however, is that this triangulation occurs as data collection proceeds (which may be distinctive to case studies) and is not the same as the triangulation that later occurs when findings are being interpreted (common for all types of empirical research).

Formal case study protocols are used to promote such triangulation during data collection, as well as to define the data collection process more generally (Yin 1994: 63–74). A protocol is not a questionnaire (posing questions to interviewees) but represents the investigator's own agenda in pursuing the line of inquiry for the case study. In other words, the "respondent" for the case study protocol is the case study investigator, addressing research questions and following a line of inquiry by having collected and triangulated a variety of evidence.

Distinguishing Evidence from Interpretation: A Case Study Database

A final desirable characteristic of case studies is to present the case study evidence separate from the investigator's interpretations of the evidence. This separation is common in laboratory and quantitative studies, in which results and data tables are presented before interpretation takes place. Unfortunately for case studies, the traditional mode of presentation has been a narrative that does not typically distinguish between evidence and interpretation—again leading to a frequent complaint that case study investigators are presenting only the evidence that supports their interpretations.

The desired remedy is to construct a case study "database" analogous to the "raw data" in a laboratory study or survey. Because the case study evidence may be quantitative or qualitative, the database may be diverse, containing in part the same kind of raw data as in a survey but also containing narrations and word tables reflecting the qualitative evidence. The database can be structured by following the line of inquiry or research questions in the case study protocol, so that evidence bearing on each question is assembled in the same place. The final case study should then draw from this database, presenting critical portions of it but citing the data as findings and not interpretation.

ANTICIPATING THE LIKELIHOOD OF HIGH-QUALITY CASE STUDIES IN REVIEWING PROPOSALS

In addition to the preceding eight characteristics of well-conducted case studies, several reminders are useful when the further challenge is to anticipate the quality of case studies on the basis of proposals. The following features are therefore specifically offered for prospective reviewers.

Investigators Seasoned in Doing Case Studies

Because case studies have traditionally been used as an exploratory method, a frequent belief is that any type of trained investigator can successfully conduct a case study. However, investigators trained in other methods are not necessarily able to conduct high-quality case studies. Similarly, whereas data collection for laboratory or survey studies can be relegated to graduate or research assistants, case study research requires constant participation—in all phases of the study, including data collection—by investigators seasoned in doing case studies. In part this is because only such investigators are likely to be rigorous in pursuing the case study protocol in the face of unexpected and diverse evidence and in part because of the need to preserve flexibility and the opportunities for discovery.

Obviously, no investigators start as seasoned case study investigators. A logical question is whether this first feature therefore represents an attempt to keep the field open only to those who already have been doing case studies. The answer is certainly not. Lesser levels of experience also are acceptable, such as a background in any qualitative or case study training in graduate school, a thesis or dissertation that might have involved such methods, or the investigator's experience as part of a multi-disciplinary team that included case study work. In the absence of any of these experiences, reviewers might still find the investigator's credentials sufficient for conducting a simple, new case study (preferably a single, not multiple-case study, to start), given that all other conditions have been met. What is to be avoided is the notion that an investigator with virtually no case study experience can credibly implement a new, multiple-case study of complex proportions. As with any other research methods, the likely result of combining inexperience with a complex research problem will, at best, be a mediocre investigation.

Review of Case Study Literature

Good proposals need to show mastery over the existing literature, from the perspectives of both substance and method. However, many case study proposals, and especially those in health services research, may underestimate the amount of earlier methodological work in case studies. Reviewers should assume that relevant case study research already exists in nearly every field. Promising proposals should be expected to reflect familiarity with the literature and even to draw analytic lessons from it.

Further, a good literature review should not be limited to the field at hand, such as health services research, but should be open to uncovering

valuable case study research in related fields, such as organizational studies or business and marketplace studies. The proposal that will lead to a high-quality case study should demonstrate a broad awareness of prior research, including the previous development of (theoretical or methodological) insights to avoid unknowing reinvention in the proposed study. The breadth of a case study literature review, along with the care with which it has been done and the quality of the insights developed, plays a major role in signaling the likely quality of the proposed case study.

Presentation of Preliminary Design, Formal Case Study Protocol, and Results of Any Pilot Testing

In case studies, as with other types of research, early discussions of the preliminary research design and the data collection plans (both reflected in a formal case study protocol, as previously mentioned) will mark the signs of a good proposal and will likely be associated with a high-quality study. Simple as this sounds, a common but mistaken position assumed in case study proposals is that planning for design and data collection must await the actual start of a study because of the need to permit flexibility and promote discovery. The discussions of these topics earlier in this article should have created a convincing argument that this position is incorrect and should not be accepted by reviewers.

On the contrary, a case study proposal may be considered even stronger if it not only contains the preliminary design and data collection plans but also shows that some pilot testing, using a pilot case study (Yin 1994: 74–76), already has been done. If nothing else, such pilot testing reduces the likelihood that the proposed case study will have to undergo a pilot phase once the actual investigation has formally begun. The potential gain, however, is much greater than this minimal savings. Pilot testing before a proposal is submitted may reveal the need for fundamental changes in a research inquiry, its design, or its data collection—and these considerations will then have been formally presented to reviewers for further feedback before any real case study has begun.

Acknowledging the Special Problem of Doing Multiple-Case Studies with Multiple Research Teams

A seemingly classic debate appears to be endemic to case study research, reflecting a broader and persistent "quantitative versus qualitative" research debate more generally (e.g., Reichardt and Rallis 1994). This debate concerns

the extent to which case study investigators are to practice more normative rules of scientific inquiry (typically branded as a "positivist" approach)—as espoused in the Reichardt and Rallis article—compared to pursuing the more intuitive insights yielded by such techniques as participant observation and grounded theory (Strauss and Corbin 1990). Nearly every rendition of this debate is cast in philosophic terms—whether human science can emulate natural science or whether "human science" might in fact be considered an oxymoron due to the overriding influences of time, space, and culture in human (but not necessarily scientific) affairs.

Reviewers of case study proposals need not become engaged in this debate unless they want to, except when they confront one critical and increasingly common situation: when the proposed case study is to consist of multiple cases conducted by multiple research teams. The situation commonly arises because of the relevance of this design to satisfying pressing policy issues in health services research. The resulting requirement that studies be completed within "reasonable" deadlines means that the lengthy period needed if a single investigator is to carry out the multiple case studies alone cannot be permitted.

Under this condition, a major assumption is that the multiple research teams need to have common orientation and training, and that they will follow a similar field protocol. Without such assurance, the value of the multiple cases would be highly questioned, because differences found among the cases could not readily distinguish between actual substantive differences and the artifactual differences due to inconsistencies from team to team. This author is unaware of any alternative for attaining such consistency without adopting the positivist posture. If this is true, then, even though the classic philosophic debate cannot be settled, the fact remains that proposed multiple-case studies are likely to be of higher quality the more they include positivist-like features such as the identification of hypotheses and rival explanations, the use of case study protocols, and the distinction between evidence and interpretation as part of the initial proposal.

CONCLUDING COMMENT

This article has addressed the problem of enhancing the quality of case studies in health services research. The main strategy has been to consider the conduct of case studies as a serious and formal research craft, not just an exploratory and informal tool (classically connoted by the term "soft research"). Investigators who want to do high-quality case studies will find

that such so-called "soft" research is actually extremely "hard" to do well (Yin 1994: 16). The guidelines in this article are intended to lessen the chore, ever so slightly, for both prospective case study investigators and reviewers of case study proposals.

ACKNOWLEDGMENTS

The author would like to acknowledge extensive discussions with Kelly Devers, Ph.D. of the Agency for Health Care Policy and Research, who both stimulated many of the ideas for the article and invited the author to participate in one of the panels in the workshop. Similarly, the other participants in the panel, Dr. Michael Quinn Patton, Dr. Gloria Bazzoli, and Dr. Martin Sanchez-Jankowski, helped through their remarks to clarify important insights. Finally, two anonymous reviewers called attention to many difficulties in reading an earlier draft, resulting in extensive revision that I hope makes key concepts clearer and the text more readable. Notwithstanding these contributions, the author alone is responsible for any errors or misinterpretations in the article.

NOTES

1. An updated and briefer version of this book is Yin (1998).
2. Of course, case studies, experiments, surveys, and histories all can be stretched so that they overlap with each other (e.g., oral histories use live, contemporary informants to study events in the recent past). The categorizations are therefore to be considered stereotypes and not limitations on investigators' inventiveness.
3. For additional information about multiple-case studies and a perspective compatible with that of this article, see Ragin in this issue.
4. Note again that the problem is not alleviated by having a few more cases but only by having a very large number, which is likely to be beyond the resources or time available. A frequent diversionary tactic is to implement an experimental or quasi-experimental design with a unit of analysis embedded within the case, if such a unit is relevant (refer to Figure 1). Such a design is fine for the embedded unit of analysis, but it still does not solve the problem at the level of the case as a whole.

REFERENCES

Bellin, S. S., H. J. Geiger, and C. D. Gibson. 1969. "Impact of Ambulatory Health Care Services on the Demand for Hospital Beds." *The New England Journal of Medicine* 280 (April): 808–12.

Clark, D. 1997. "What Is Qualitative Research, and What Can It Contribute to Palliative Care?" *Palliative Medicine* 11 (March): 159–66.

Ginsburg, P. 1996. "The RWJF Community Snapshots Study: Introduction and Overview." *Health Affairs* 15 (summer): 8–20.

Mays, N., and C. Pope, eds. 1996. *Qualitative Research in Health Care.* London: BMJ Publishing Group.

Nelkin, D. 1973. *Methadone Maintenance: A Technological Fix.* New York: George Braziller.

Neustadt, R. E., and H. Fineberg. 1983. *The Epidemic That Never Was: Policy-Making and the Swine Flu Affair.* New York: Vintage Books.

Patton, M. Q. 1999. "Enhancing the Quality and Credibility of Qualitative Analysis." *Health Services Research* 34 (5, Part II): 1189–1208.

———. 1990. *Qualitative Evaluation and Research Methods, 2d ed.* Thousand Oaks, CA: Sage.

Ragin, C. C. "The Distinctiveness of Case-oriented Research." *Health Services Research* 34 (5, Part II): 1137–52.

Ragin, C. C., and H. S. Becker, eds. 1992. *What Is a Case? Exploring the Foundations of Social Inquiry.* New York: Cambridge University Press.

Redman, E. 1973. *The Dance of Legislation.* New York: Simon & Schuster.

Reichardt, C. S., and S. F. Rallis, eds. 1994. "The Qualitative-Quantitative Debate: New Perspectives." *New Directions for Program Evaluation* 61 (spring).

Sofaer, S. "Qualitative Methods: What Are They and Why Use Them?" *Health Services Research* 34 (5, Part II): 1101–18.

Srole, L., T. S. Langner, S. T. Michael, M. K. Opler, and T. A. Rennie. 1962. *Mental Health in the Metropolis.* New York: McGraw-Hill.

Strauss, A., and J. Corbin. 1990. *Basics of Qualitative Research: Grounded Theory Procedures and Techniques.* Thousand Oaks, CA: Sage.

U.S. Office of Technology Assessment. 1979–1981. *Case Studies of Medical Technologies.* Washington, DC: OTA. (17 cases published during the three-year period.)

Wholey, J. 1979. *Evaluation: Performance and Promise.* Washington, DC: The Urban Institute.

Yin, R. K. forthcoming. "Rival Explanations as an Alternative to Reforms as 'Experiments.'" In *Validity and Social Experimentation: Donald Campbell's Legacy,* edited by L. Bickman. Thousand Oaks, CA: Sage.

———. 1998. "The Abridged Version of Case Study Research." In *Handbook of Applied Social Research,* edited by L. Bickman and D. J. Rog, pp. 229–59. Thousand Oaks, CA: Sage.

———. 1994. *Case Study Research: Design and Methods, 2d ed.* Thousand Oaks, CA: Sage.

Yin, R. K., and D. Yates. 1975. *Street-Level Governments: Assessing Decentralization and Urban Services.* Lexington, MA: D. C. Heath.