Distinctions among Different Types of Generalizing in Information Systems Research

Allen S. Lee
Virginia Commonwealth University
1015 Floyd Avenue
Richmond, Virginia 23284-4000
USA
AllenSLee@alum.mit.edu
PHONE: (804) 827-0366
FAX: (804) 828-3199

Richard Baskerville
Georgia State University
35 Broad Street
Atlanta, Georgia 30303
USA
baskerville@acm.org
PHONE: (404) 651-3880
FAX: (404) 651-3842

The authors presented an earlier version of this paper at the Working Conference of Working Group 8.2 of the International Federation for Information Processing, which was held in St. Louis, Missouri, August 20-22, 1999. The current version accounts for comments and feedback received at the conference.

Distinctions among Different Types of Generalizing in Information Systems Research

Abstract

It is incorrect and even harmful that many information systems researchers typically criticize their own intensive (qualitative, interpretive, critical, and case) research as lacking "generalizability." We untangle and distinguish the numerous concepts now confounded in the single term "generalizability," which are *generality*, *generalization*, *generalize*, *general*, and *generalizing*. These clarified terms allow us to identify four distinct forms of generalizing (everyday inductive generalizing, everyday deductive generalizing, scientific inductive generalizing, and scientific deductive generalizing), each of which we illustrate with an information-systems-related example. With these clarified terms, we show how information systems researchers who perform intensive research can properly lay claim to generality for their research.

Distinctions among Different Types of Generalizing in Information Systems Research

For researchers in information systems (IS), the concept of "generalizability" has been developing and maturing with the growing acceptance of intensive research approaches (e.g., qualitative, interpretive, critical, and case research[1]). Often misconstrued to be a property of statistically based research alone, generalizability has been gaining recognition as an achievable ideal in intensive research as well (Lee, 1989). Still, it remains common for intensive researchers to flagellate themselves in their published articles' "discussion" sections for the lack of generalizability of their findings. There they typically blame this supposed failure on their having examined "only" a single case, or "only" three technologies, or "only" two organizations, or "only" one point in time, and so forth. We believe that such self-flagellation is not necessary (i.e., the completed intensive research can indeed claim "generalizability" if it is properly performed and presented). Such unwarranted self-criticism can even be harmful (i.e., it incorrectly elevates large-sample statistical research above all other forms of scholarly inquiry). The purpose of this study is to clarify the different processes of generalizing so that scientific researchers in IS can better achieve and securely claim "generalizability" (or, as we will rename it, "generality") in their research. We will illustrate these processes with some examples of how authors of specific published articles in IS research can better describe the "generalizability" of their studies.

The first section of the paper after this introduction will identify and distinguish different terms now confounded in the single term, "generalizability." The terms that we will unconfound and extract from "generalizability" are *generality*, *generalization*, *generalize*, *general*, *generalizing*, and even (after we define it) *generalizability* itself. In the section after that, we will define four types of generalizing and classify them according to the dimensions of "reasoning process" (inductive *vs.* deductive generalizing) and "context" (inquiry in everyday life *vs*. inquiry in scientific research). The same section will illustrate the four types with

[1] We take this characterization of "intensive research" from M. Lynne Markus and Allen S. Lee's call for papers for a special issue of *MIS Quarterly* on intensive research. They, in turn, took the term "intensive research" from Karl Weick and used it to refer to the diversity of forms of empirical information-systems research falling outside of the quantitative and positivist genre, including qualitative positivist (and non-positivist) research, interpretive research, critical social theory research, and case study research.

empirical examples from the published literature. In the third section, we discuss an appropriate way in which IS researchers may indeed lay claim to generality for their research. Then, in the last section, we will bring out the ramifications of the four distinct forms of generalizing for current and future research practices in IS. There we will indict and dismiss the dysfunctionality of self-flagellation for the often imagined sin of "lack of generalizability" in intensive research, as well as proclaim our emancipation to a research environment with a better (or "generalized," as it were) conception of "generalizability."

1. UNCONFOUNDING AND RENAMING "GENERALIZABILITY"

In the way that scientific researchers now use the term "generalizability," it regards the extent to which a scientific researcher's theory does or does not apply to empirical referents (i.e., real-world situations) apart from the one that the researcher examined in his or her study. A case researcher, for example, might offer the self criticism (in the discussion section of her published case study) that the theory she developed "lacks generalizability" because she based it on observations of only one information technology in only one organization. However, whether or not we would agree with her self criticism, we believe that she means *generality*, not *generalizability*.

In this study, we use the term *generalizability* more specifically. We distinguish it from *generality*, *generalization*, *generalize*, *general*, and *generalizing*. We define *generalizability* to refer to a theory's potential (as the suffix in *generaliz<u>ability</u>* signifies) to come to possess the quality of *generality*. *Generality* refers to the range of phenomena across which the theory has been demonstrated to hold. Here, *generality* is the outcome or product of the process of *generalizing* a theory that, initially, was able to be generalized (i.e., *generalizable*). To *generalize* is to engage in the process of *generalizing*. A theory of perfect *generality* would apply to the entire class of empirical referents that it purports to explain; a theory with less *generality* would apply only to a subset of this class.[2] In other words, we identify *generalizability* as a property of a theory at the beginning of an empirical investigation, and

[2] We are aware of no published empirical study in the field of information systems that has ever offered *explicit* statements about what constitutes the entire class of empirical referents that its theory purports to explain. Even our own publications (i.e., the research that the authors of this paper have published) do not do this. It would be fair and justifiable to say that researchers generally imply what the entire class is (e.g., it could be "all managers"

generality as a property of a theory at the end of the investigation (where the results of the investigation are favorable). The resulting *generalization*, in this scheme, would be the theory in the form in which it emerges from the empirical investigation. Finally, a theory possessing *generality* can also be described with the adjective *general*.

Using the terms *generality*, *generalization*, *generalize*, *general*, *generalizing*, and *generalizable*, we now proceed to identify four types of generalizing.

2. FOUR TYPES OF GENERALIZING

To help clarify the different processes of generalizing so that scientific researchers in IS can better achieve and claim generality in their research, we offer a framework that draws attention to four types of generalizing.

First, we recognize that the act of generalizing is not something done only by scientific researchers, but also something done by everyday people in everyday life.[3] Consultants and managers, for instance, can (and, arguably, must) generalize from just one or two experiences. We therefore define one dimension of generalizing as referring to the context of inquiry. This context could be the inquiry of a person in everyday life (such as a consultant or a manager) or alternatively the inquiry of a person in scientific research (such as an IS researcher). In the context of inquiry in everyday life, the generality that a person associates with a belief is largely determined by the social traditions of the person's group. In contrast, in the context of inquiry in scientific research, the influence of social traditions of the scientific group in establishing generality are strongly mediated by rigorous notions of evidence, logic, and methodology in "scientific" thinking. Second, we recognize two different conceptions of the process of generalizing. These are inductive generalizing (a reasoning process that begins with observations and subsequently uses them as the basis on which to build a theory) and deductive generalizing (a reasoning process that begins with a theory and subsequently processes or tests it against observations). In **Table 1**, we use these two dimensions to identify four types of generalizing. We will now examine each of the four types in greater detail. Each type will be

---

or "all corporations"). Still, we believe that, as a matter of good methodology, it would be a good idea for empirical studies to adopt this practice. We flesh out this recommendation in this study's discussion section.

[3] Furthermore, any distinctions between everyday generalizing (which occurs in what Schutz, 1973, calls the "natural attitude" of everyday life) and scientific generalizing (which occurs in what Schutz calls the "scientific attitude" of the observing researcher) will also be useful for better distinguishing and sharpening what we will mean

illustrated by means of an IS-related example. We chose these examples not only because they illustrate the principles, but also because there are published details that illuminate the generalization process in each example.

FOUR TYPES OF GENERALIZING

		CONTEXT	
		Inquiry in Everyday Life	**Inquiry in Scientific Research**
REASONING PROCESS	**Induction**	*everyday inductive generalizing* "typification"	*scientific inductive generalizing* "generalizability"
	Deduction	*everyday deductive generalizing* "learning"	*scientific deductive generalizing* "falsification"

2.1 Everyday Inductive Generalizing

The sociologist and phenomenologist Schutz (1973) offers an explanation of how everyday people in everyday life make generalizations. "All projects of my forthcoming acts," as an everyday person in everyday life, "are based upon my knowledge at hand at the time of projecting," Schutz explains (p. 20). "To this knowledge belongs my experience of previously performed acts which are typically similar to the projected one. … The first action $\boldsymbol{A'}$ started within a set of circumstances $\boldsymbol{C'}$ and indeed brought about the state of affairs $\boldsymbol{S'}$; the repeated action $\boldsymbol{A''}$ starts in a set of circumstances $\boldsymbol{C''}$ and is expected to bring about the state of affairs $\boldsymbol{S''}$." Regarding a person's general conception of a type of action, Schutz acknowledges that no

by the latter.

two real-world instances or instantiations of it (e.g., ***A′*** and ***A″***) are ever exactly the same, which leads to the following point regarding how a person in everyday life makes a generalization:

> Yet exactly those features which make them unique and irretrievable in the strict sense are – to my common-sense thinking – eliminated as being irrelevant for my purpose at hand. When making the idealization of 'I-can-do-it-again' I am merely interested in the typicality of ***A***, ***C***, and ***S***, all of them without primes. The construction consists, figuratively speaking, in the suppression of the primes as being irrelevant, and this, incidentally is characteristic of typifications of all kinds.

What Schutz calls a "typification" is further discussed by Berger and Luckmann in their classic book *The Social Construction of Reality* (1966). In one person's interactions with another person, the former could "apprehend the latter as 'a man,' 'a European,' 'a buyer,' 'a jovial type,' and so on" (Berger and Luckmann, p. 31), where "these typifications ongoingly affect" how the former interacts with the latter. Typifications serve the purpose of allowing everyday people to negotiate their interactions with one another, as well as with the physical world around them. Returning to Schutz's example, we denote the process of generalizing in everyday life as follows:

$$\boldsymbol{A', A'', A''', \ldots} \implies \boldsymbol{A}$$

Another example, again referring to Schutz's typifications, would be:

$$\{\boldsymbol{A', C', S'}\}, \{\boldsymbol{A'', C'', S''}\}, \{\boldsymbol{A''', C''', S'''}\}, \ldots \implies \{\boldsymbol{A, C, S}\}$$

The reasoning process that produces such typifications is what we call "everyday inductive generalizing." We describe this reasoning process as "everyday" because it refers to inquiry in everyday life, rather than inquiry in scientific research. We identify this reasoning process to be a form of "inductive generalizing" because it begins with observations and subsequently uses them as the basis on which to generalize and then construct the typification.

Example of everyday inductive generalizing in IS

For an example, we will focus on one particularly interesting typification constructed by the professionals who work in the everyday world of IS development. It is what they call the "death march project."

They have come to see the "death march project" in the following way. It involves IS projects where signs of failure are apparent, but all participants in the project nonetheless proceed to play their parts, as if there were nothing wrong. "Death march projects" have come to be known, at least to those in the everyday world of IS development, by certain notable signs. One sign is that the risk of project failure appears greater than 50 percent. Other frequent signs of a "death march project" include irrational managerial compensation, the drowning of staff with a complete, sudden conversion to a new silver-bullet methodology, and the driving away of staff through cancellation of vacations and weekends (Yourdon 1997). Through everyday inductive generalizing across their observations of such IS development settings, software engineering practitioners have constructed the "death march project" typification. The sorts of projects now seen as "death march" projects are not new, having existed since the 1960's (Yourdon 1997); however, the typifying or generalizing of these observations into the shared, everyday concept of "death march" occurred only in the 1990's.

The process by which the "death march" typification results from the everyday inductive generalizing performed by a group of developers is illustrated in the February 1997 issue of *American Programmer*, which is dedicated to the topic of death march projects. In this issue, corporate IS practitioners and IS consultants speak from their experiences with death march projects. Each death march project in their experience constitutes a discrete phenomenon that they have experienced, where $\boldsymbol{P'}$ could be the SMS/800 nine hour project (Oxley 1997) and $\boldsymbol{P''}$ could be the billing and accounts receivable system (Roberts 1997). Across these discrete experiences, they have generalized to $\boldsymbol{P}$, the death march project:

$$\boldsymbol{P', P'', P'''} ===> \boldsymbol{P}$$

$\boldsymbol{P}$, the "death march project" typification, is the general case. The process of generalizing commences by noting characteristics shared across the observed cases $\boldsymbol{P', P'', P'''}$, *etc.* Examples of these shared characteristics are certain irrational management actions noted above. The existence of the "death march project" was never "hypothesized" or "tested" in the scientific sense of these terms. Rather, this general case is a consequence of the generalizing by everyday IS developers across certain unpleasant projects. The general case is significant to these developers because of its instrumentality. The developers use the concept as a normative model to suggest possible actions when they encounter death march projects in the future. Though not tested for the status of scientific truth, the "death march" typification acquires the status of everyday truth "in so far as [it] helps us to get into satisfactory relation with other parts of our experience" (James, 1975, p. 35).

2.2 Everyday Deductive Generalizing

Not all generalizations in everyday life are the "outputs" of a reasoning process for which observations of particular instances are the "inputs." In other situations, an already existing generalization is itself the "input" to a reasoning process that applies it, where the "output" is the result that the generalization suited, or failed to suit, the application.

Argyris and Schön (1978) have already recognized that everyday people hold such theories and test them for suitability in new applications. Argryris and Schön distinguish between a person's "espoused theory" and the same person's "theory-in-use." A person's espoused theory is the explanation that this person would voice to explain her behavior. A person's theory-in-use is the theory that actually governs this person's behavior. The theory that a person espouses is not always the same as the person's theory-in-use; indeed, a person might not even be aware of her theory-in-use. In the immediately following discussion, we will examine (for the sake of simplicity) only espoused theories.

When a person's actions repeatedly elicit reactions from other people or her environment that surprise her – that is, when her espoused theory apparently does not work – the person can sense an error or deficiency in her espoused theory. If the person indeed becomes conscious of the error or deficiency, she can then proceed to question her espoused theory and then to revise it.

The everyday reasoning process in this behavior is not inductive, but deductive; it begins not with the observations of particulars, but with an already existing generalization (the person's espoused theory). The person applies this generalization to a set of particulars she observes in the empirical setting in which she acts (rather than derive the generalization from her observations of these particulars), leading to expectations of how other people and her environment should (if the generalization is right) react to her. In this reasoning process, two results are possible. In one, the person becomes surprised when other people or her environment reacts to her in ways that she would not expect. This would set the stage for her to learn that her espoused theory requires revision. In the other result, there is nothing in the environment's or other people's reactions to her that she would find surprising or unexpected, whereupon her espoused theory would survive and become, in a sense, further entrenched or stronger as a generalization. We refer to this reasoning process as "everyday deductive generalizing."

Argyris and Schön were hardly the first scholars to recognize that deductive generalizing occurs in everyday life. For instance, this form of generalizing is the hallmark of the philosophy of pragmatism as found in the respective works of William James, John Dewey, and Charles Sanders Pierce.

Example of everyday deductive generalizing in IS

For an example, we will focus on the generalizing in which some people engaged in an IS action research project.

Action research is a research method that aims to solve immediate practical problems while expanding scientific knowledge. Based on collaboration between researchers and research subjects, action research is a cyclical process that builds learning about change in a given social system (Hult and Lennung, 1980). Unlike *laboratory* experiments that isolate research subjects from the real world, action research involves *intervention* experiments in which the researchers, along with the research subjects, apply a stimulus or other change strategy to the real-world context in which the research subjects live or work. In action research, intervention experiments operate on problems or questions that the practitioners (to whom we also refer interchangeably as research subjects and everyday people) themselves perceive within the context of their own particular empirical setting. *Participatory* action research is distinguished by the additional characteristic of involvement of, first, practitioners not only as subjects but also as co-researchers and, second, researchers not only as scientists, but also as subjects. "It is based on the Lewinian proposition that causal inferences about the behavior of human beings are more likely to be valid and enactable when the human beings in question participate in building and testing them" (Argyris and Schön, 1991, p. 86). Because of the deep collaboration between scientific researchers and everyday professionals, participatory action research involves both scientific and everyday deductive generalizing. Consequently, examples of this form of research can illustrate the two categories of generalizing. However, in this portion of our study (section 2.2), we will focus on how the everyday professionals themselves engage in deductive generalizing. (We will defer our discussion on how the scientific researchers engage in scientific deductive generalizing until section 2.4.)

One example of an action research project in IS appears in Baskerville and Stage's study (1996) on risk management for prototyping. We may summarize their study as follows. It addresses one of the main problems in the overall practice of prototyping that, at least at the time of the study, had not yet been resolved. The practical problem was the difficulty simply in controlling the scope and unfolding development of prototyping projects. In general, this problem can severely limit the range of IS development projects in which prototyping can be used effectively. In Baskerville and Stage's study, the action research project developed and validated a new approach that uses an explicit risk mitigation model in the IS development process, one that focuses the collaborative action research team's attention on the consequences and priorities inherent in the prototyping situation. The study established that prototyping projects could be controlled if appropriate risk resolution strategies were put into effect prior to any breakdown in the prototyping process.

Baskerville and Stage's study illustrates how collaborative action team members engage in everyday deductive generalizing. In the example that follows, the generalizing pertains to the team members' original conception of a prototyping risk factor that they called "user alienation."

Specifically, while in the process of applying risk analysis as a tool to help administer their prototyping project, the collaborative action research team members stated certain expectations during an initial analysis of the risk factors that they perceived in the problem setting of their practical prototyping. Using the Argyris and Schön terminology that we introduced above, we can say that the collaborative action research team members came to "espouse" a theory that named twelve risk factors. At the same time, because they actually applied and followed their "espoused theory," we may state that, in this case, the "espoused theory" coincided with their "theory-in-use."

The collaborative action research team members stated one of the twelve risk factors as, "Users will not understand what we are doing" (Baskerville and Stage 1996, p. 493) and they stated the consequence of this factor as, "The users become alienated." They ranked this factor's risk level as moderate.

However, the participants' understanding of risk factors led them to sense and identify a reaction from the first prototyping cycle as being surprising. The participants noted that the mentioned risk factor carried not just (as expected) one, but actually two separate consequences. The first of the two separate consequences was enduring and serious: "The users do not know what product they will receive" (p. 495). The second was transient and trivial: "The users do not understand their role in the development process." The need to differentiate the original single consequence into two became obvious during the first learning cycle as the users gained experience with prototyping and adapted enthusiastically to their role in the development process (Baskerville and Stage 1991, p. 12).

To recapitulate, we note that the intervention experiment was not inductive, but deductive; the intervention experiment began not with the observations of particulars, but with a generalization about risk factors that the collaborative research action team members had previously formulated. As required in the invention phase of action research, the researchers were not observing as detached scientists, but instead were participating in the role of everyday people themselves who are acting and thinking in the real world. In their intervention, all the team members applied their already existing generalization to a set of particulars that they observed in the IS development context in which they sought to act. In the results from their intervention, they noticed the major surprise in which there was not just one, but two different, consequences to one of their theorized risk factors, whereupon they learned that their understanding needed to be changed. We refer to this reasoning process as an instance of everyday deductive generalizing. Moreover, in this particular instance, the conclusion was not to claim greater generality for their initial understanding, but rather, to establish that it lacked generality. The everyday people *learned* that they needed to change their already existing generalization.

We may summarize the above explanation symbolically as follows. Using the same Schutz-based terminology of {***A*** , ***C***, ***S***} as we used in section 2.1, we can say that the everyday people learned that their already existing generalization, {***A*** , ***C***, ***S***}, lacked generality.

Specifically, in the experience of the first prototyping cycle, the instantiation of their generalization took the form, $\{\boldsymbol{A'}, \boldsymbol{C'}, \boldsymbol{S'}\}$. However, rather than observing the expected state of affairs $\boldsymbol{S'}$ as resulting from engaging their specific action $\boldsymbol{A'}$ in the first prototyping cycle's circumstances $\boldsymbol{C'}$, they instead encountered $\neg\boldsymbol{S'}$ (i.e., "not S prime"). This instance indicated to the everyday people doing the prototyping that their already existing generalization $\{\boldsymbol{A}, \boldsymbol{C}, \boldsymbol{S}\}$ lacked generality, whereupon they realized that $\boldsymbol{S}$ (indicating the state of affairs in which there is one consequence to the action $\boldsymbol{A}$) did not fit and that they needed to replace it with $\boldsymbol{Q}$ (indicating the state of affairs in which there are two consequences to the action $\boldsymbol{A}$) in their generalization. Thus the resulting generalization was $\{\boldsymbol{A}, \boldsymbol{C}, \boldsymbol{Q}\}$.

2.3 Scientific Inductive Generalizing

We observe that, among many (and sometimes, it seems like most) IS researchers, there is the (mis)conception that the greater the sample size or number of observations ($\boldsymbol{O'}, \boldsymbol{O''}, \boldsymbol{O'''}$, etc.), then the more "generalizable" (or, in the terminology of our study, the more general) the resulting theory $\boldsymbol{T}$ is. We use the term "scientific inductive generalizing" to refer to the process by which social scientific researchers would begin with observations (such as "n" sample points) and end up with a theory. Using the Schutz-inspired terminology that we introduced earlier in this paper, we may represent the proposition, "inductive inference leads to a valid theory," as follows:

$$\boldsymbol{O'}, \boldsymbol{O''}, \boldsymbol{O'''}, \boldsymbol{O''''}, \ldots, \boldsymbol{O''''^{\cdots}} \Longrightarrow \boldsymbol{T}$$

Whereas this statement is intuitively appealing, it is fatally flawed. The theory, $\boldsymbol{T}$, is more complex than the suppressed-prime typification, $\boldsymbol{O}$, resulting from everyday inductive generalizing. $\boldsymbol{T}$ can additionally refer to what is not observable but must be inferred, such as causal mechanisms, certain relationships among factors, and even certain factors themselves. Theories pertain to, among other things, "*unobservable* entities, such as molecules and atoms, electrons and protons, chromosomes and genes" (Copi, 1986, p. 483). Consider that even the natural sciences, such as evolution, astronomy, and geology, refer to phenomena that occurred long ago and whose existence is no longer accessible to direct observation today, but only to theorizing. Conversely, consider that if all the phenomena of interest to a researcher were readily observable, there would be no need for the researcher to theorize about these phenomena in the first place. Because there is no logic by which to transform statements about what is observable (such as $\boldsymbol{O'}$, $\boldsymbol{O''}$, and $\boldsymbol{O'''}$) into statements about what is unobservable (i.e., a theory ***T***), any attempt to implement the statement "inductive inference leads to a valid theory" would be infeasible.[4]

Lee (1999) reviews, from the philosophy of science, another argument recognizing that, while induction is fine for suggesting the formulation of a theory, induction offers no help in empirically testing or validating a theory. In particular, consider whether the proposition, "inductive inference leads to a valid theory," is itself true. Lee (p. 14) credits Popper (1968, p. 29) for this argument, who states it in the following way:

> That inconsistencies may easily arise in connection with the principle of induction should have been clear from the work of [the philosopher David] Hume… Thus if we try to regard its truth [i.e, the truth of the principle of induction] as known from experience, then the very same problems which occasioned its introduction will arise all over again. To justify it, we should have to employ inductive inferences; and to justify these we should have to assume an inductive principle of a higher order; and so on. Thus the attempt to base the principle of induction on experience breaks down, since it must lead to an infinite regress.

In other words, to justify the proposition "inductive inference leads to a valid theory" as being true, one would marshal "n" examples where, in each example, one would observe inductive inference leading to a valid theory. (For reasons that will shortly become clear, we will refer to this as a "first order" inductive inference.) However, this would create an obvious problem: one

---

[4] Indeed, if the observations $\boldsymbol{O'}$, $\boldsymbol{O''}$, $\boldsymbol{O'''}$, $\boldsymbol{O''''}$, … , $\boldsymbol{O''''^{\cdots}}$ can be generalized to anything, they can be generalized to ***O***, not ***T***:

$$\boldsymbol{O'}, \boldsymbol{O''}, \boldsymbol{O'''}, \boldsymbol{O''''}, \ldots, \boldsymbol{O''''^{\cdots}} \Longrightarrow \boldsymbol{O}$$

For instance, the field of statistical inference may properly generalize from the observations of incomes of individuals randomly selected from a population to an estimated mean income for the population. Indeed, such induction may properly estimate the mean incomes of different subgroups in the population, such as men, women, whites, and people of color. However, these observations of incomes of individuals may not and cannot, *by themselves*, enable the derivation or induction of theoretical statements about the mechanisms causing women's incomes to be lower than men's or the incomes of people of color to be lower than the incomes of whites.

would be applying inductive inference to justify itself. To justify this first-order application as valid, one would then marshal "n" examples of such applications. (This would be a "second order" inductive inference.) However, this would again incur a problem: one would again be applying inductive inference to justify itself. To justify this second-order application as valid, one would then marshal "n" examples of such applications. (This would be the "third order" inductive inference.) However, this would again incur a problem: one would again be applying inductive inference to justify itself. To justify this third-order application as valid, one would then marshal "n" examples of such applications. And so forth. In short, this argument shows that the purported methodological principle, "inductive inference leads to valid theories," is itself not empirically justifiable because any attempt to establish its own empirical validity would ultimately depend on induction itself, hence leading to the "infinite regress" that Popper mentions.

The final argument we will cite against inductive inference is that the same set of observations can be consistent with two (or more) incompatible theories at the same time. To illustrate this argument, consider that, with every passing day, we can increase the number of observations in which the sun rises and the sun sets, where this growing number of observations would support *both* the theory that the earth is at the center of the solar system and the contrary theory that the sun is at the center of the solar system. Hence, if inductive inference were a legitimate means of justifying a theory, it would be justifying *both* of these contradictory theories as true.

Lee (p. 15) notes that, interestingly, even the discipline of statistical inference has distanced itself from the notion of induction, where statistical inference recognizes that a larger sample size, "n," does not increase the probability that a statistically inferred proposition (such as a confidence interval) is true, but instead serves to enhance the "level of confidence" or "statistical significance," which is an attribute of the researcher's investigation, not an attribute of the proposition.

As we will explain in the next section (2.4), deductive generalizing is the viable alternative to inductive generalizing in scientific research.

Examples of scientific inductive generalizing in IS

In this section, we will offer some examples that describe how IS researchers use an inductive framework not only to help formulate a theory (a use that we see as noncontroversial), but also to lament (in our view, incorrectly) the supposed lack of "generalizability" (generality) of their own empirical studies.

One example involves the qualitative research approach of “grounded theory,” which does just what its name implies: it attempts to ground theory on observation. “A grounded theory is one that is inductively derived from the study of the phenomenon it represents” (Strauss and Corbin, 1990, p. 23). As such, the grounded theory approach well exemplifies how induction can help in the formulation[5] of a theory. Orlikowski (1993) conducted a grounded-theory study of the adoption and use of CASE tools in two organizations; she states (p. 310):

> While the findings of this grounded theory are detailed and particularistic, a more general explanation can also be produced from the results… Yin (1989) refers to this technique as “analytic generalization”… Here the generalization is of theoretical concepts and patterns [from the particulars observed in a set of empirical circumstances]. This generalization is further expanded in this paper by combining the inductive concepts generated by the field study with insights from existing formal theory… The outcome is a general conceptualization of the organizational changes associated with adopting and using CASE tools that should both contribute to our research knowledge and inform IS practice…
>
> …drawing on the rich data of two organizations’ experiences, the paper generates a grounded understanding of the changes associated with implementing CASE tools in systems development.

Another example of scientific inductive generalizing involves ethnography. Schultze (forthcoming) writes about her own ethnographic study:

> …when it comes to analyzing the material collected during the fieldwork, the primary activity involves reading the fieldnotes over and over again in order to categorize events and to inductively construct themes. This typically implies manipulating the fieldnotes by cutting them apart and reassembling them into coherent patterns that apparently ‘emerge’ from the data themselves…

[5] We emphasize that inductive inference is perfectly satisfactory for use in suggesting or formulating a theory, where this is different from the empirical testing or validating of a theory.

As Orlikowski and Schultze exemplify immediately above, induction is certainly useful for suggesting and formulating a theory. However, as we have previously explained, inductive inference is fatally flawed as a means for empirically testing or validating a theory. Still, many IS researchers hold themselves to it as the standard for establishing whether or not their research possesses what they call "generalizability" (or what we call generality). And almost always, they apply this standard (which they do not know to be incorrect) so as to conclude (incorrectly) that their research lacks "generalizability" (generality). We find instances of this when we look at IS researchers who conduct case studies. They regard the single case or small number of cases that they examined to dampen the "generalizability" (generality) of their results. Applying the same logic, they say that "generalizability" (generality) can result only by increasing "n."

> From the point of view of theory development, while case studies provide useful anecdotal information, the generalizability from one specific instance to another is often limited. (Albers, Agarwal, and Tanniru, 1994, p. 94.)

> First and foremost, it should be reaffirmed that the single case research strategy employed here only allows generalizability to a research model, which in turn needs to be tested under a multiple case study design or by other field methods. (Brown 1997, pp. 90.)

> From the evidence of the two cases, it was not possible to identify any generalisable strategies for overcoming constraints but the particular solutions developed appeared to reflect the developers' local conditions and their knowledge, intuition, and experience. This would suggest that rather than giving a set of generalized guidelines for improving user involvement (as is common in the literature), the emphasis might be better placed on supporting developers' ingenuity and improvisation and on developing their social skills to enable them to overcome the constraints on involvement. (Nandhakumar and Jones 1997, p. 84.)

> Because they are drawn from a study of two organizations, these results should not be generalized to other contexts. Each context is different, so we should expect different contextual elements to interact with technical initiatives to produce different consequences. The findings should not even be extended to other settings where GIS, or even Arc/Info, is implemented. What is true for GIS in the two local county governments studied may be untrue for GIS in other governmental units or in private enterprises. (Robey and Sahay 1996, p. 108.)

> In particular, in-depth analysis of extensive data from only one organization reduces generalizability, but increases correspondence to reality. (Hidding 1998, p. 311.)

> The study has a number limitations that need to be considered in making any conclusions. First, the single case site limits the generalizability of results. The purpose of the study was not to provide generalizability of empirical results to other firms, rather the purpose was to "expand and generalize theories" (Yin 1984). (Jarvenpaa and Leidner 1997, p. 408.)

> Conducting additional case studies will provide instances of the various learning/outcome combinations, and we encourage such research. On the other hand, case studies alone will not result in validity or generalizability. Toward that end, a more fruitful approach might be to compare the development processes for similar systems in different organizations, or two or more systems being developed in a single organization. (Stein and Vandenbosch, 1996.)

The irony is that these IS researchers' method-in-use (with respect to generalizing) – the notion that a larger sample size or a greater number of observations leads to greater "generalizability" (generality) – is itself invalid. This, in turn, means that these researchers should consider themselves to be freed from what it requires. Indeed, given the invalidity of this notion, we can take the position that the authors of the reported cases studies can still assert the outcome of "generalizability" (generality) from their research, which we explain in the next section, 2.4. As for IS researchers who flagellate themselves for a supposed lack of "generalizability" (generality) in their case studies, we believe that they are not only doing this unnecessarily, but also harming the overall reputation of the scientific discipline of IS. This reputation is diminished in two ways: first, by reinforcing and perpetuating an incorrect notion of science (i.e., empirical validity is justified through induction) and, second, demeaning (and thereby discouraging the dissemination of) their own research findings whenever they draw public attention to a flaw that does not really exist.

2.4 Scientific Deductive Generalizing

Lee (1999) also reviews arguments that explain how *deductive* inference operates in science. He notes that, in deductive inference, (1) the researcher begins with an already formulated theory (which could be her own or one that she derives from the literature in her school of thought), (2) the researcher applies the statements comprising this theory to the facts describing a specific situation (such as the initial conditions in an experiment), allowing the deductions of predictions of what the researcher subsequently ought to observe or not to observe *if the theory is true*, upon the application of the experimental stimulus or treatment, and (3) the researcher, in comparing what the theory predicts in this specific situation and what she actually observes in this specific situation, ends up testing the theory deductively and indirectly. If the prediction turns out to be false, then the theory from which it followed deductively would be false too. On the other hand, if the prediction turns out to be true, then "the truth of the theory (from which the prediction originated) is not proven, but is only 'corroborated,' 'supported,' or 'confirmed' in the instance of this single test" because other new instances would open up the same theory to yet new opportunities for its falsification. Hence, in the deductive logic of empirical science, a research may never assert a theory to be correct or true, but may assert, at best, only that it is consistent with the facts. The ever-present possibility for a future empirical test to overturn it means that the theory must forever maintain the status of being falsifiable. Lee states (p. 15), "the widespread characterization of theories, even in the social sciences, as falsifiable, testable, refutable or disconfirmable [is] an indication of the widespread extent to which the deductive testing of theories is practiced."

Using the Schutz-based terminology, we may express the above argument as follows. Instead of having $\{A', C', S'\}$, where A' refers to an instantiated action, C' refers to an instantiated set of circumstances, and S' refers to an instantiated state of affairs that A' engenders in C', we have $\{T, C', S_P', S_O'\}$, where T refers to the scientist's theory, C' refers to the specific empirical conditions (the "initial conditions") instantiated in an experiment or field setting, S_P' refers to the specific state of affairs expected and predicted (hence the subscript $_P$ in S_P') to follow in the experiment if the theory is true, and S_O' refers to the specific state of affairs actually observed (hence the subscript $_O$ in S_O') to follow in the experiment or the field setting. The theory T is falsified if S_P' and S_O' differ. The theory T is corroborated or confirmed if S_O' matches S_P', but remains falsifiable in future experiments involving the new circumstances, C'', and then C''', and then C'''', and so on.

For scientific deductive generalizing, would a larger number of observations consistent with the theory being studied increase the generality attributable to the theory? To answer this question, we note that one can increase the number of observations consistent with Newtonian physics, but this hardly makes Newtonian physics more general (indeed, it is false and has been superseded by Einsteinian physics). Likewise, to increase the number of observations consistent with Ptolemaic astronomy is indeed possible, but hardly makes the theory of an earth-centered solar system more general! Therefore, in scientific deductive generalizing, simply increasing the number of supportive observations alone (such as increasing the "sample size") *does not render the theory any more general or true.*[6] This also means that the only observations that a deductive scientific researcher may regard as contributing useful information in gauging whether a theory is true or false are those that contradict the theory; such empirical evidence would, at best, throw a theory's validity and hence generality in doubt and, at worst, result in the researcher's conclusion that the theory is false and hence possesses no generality.

Examples of scientific deductive generalizing in IS

For an illustration of scientific deductive generalizing, we refer once again to the Baskerville and Stage action research study (1996) that we introduced in section 2.2, as well as to a case study by Paré and Elam (1988).

The Baskerville and Stage example involves the detection of an important delusory element in the researchers' risk analysis. This element materialized when a risk factor, though highly expected, failed to develop as time progressed. The risk factor, "The technical environment is unreliable" (p. 493), was included because the prototyping project was forced to use a recent release of a database software package known sometimes to corrupt database files.

The software performed flawlessly in the early stages of the prototype construction. Based on this experience, the members of the collaborative action research team downgraded the probability of the risk factor in question. In this scenario, we can see that the researchers posited a hypothesis ("this software is unreliable") and conducted an experiment (prototype construction) that failed to support the hypothesis – hence (deceptively) encouraging them to attribute less generality to it. However, as the prototype development continued, its functionality and complexity rose. With this intensified usage, many previously unused features

[6] As we will continue to elaborate in Section 3, a scientific researcher does not enhance generality by simply increasing the number of supportive observations, but rather, by varying the range of different empirical circumstances in which the observations, consistent with the theory's predicts, are made.

of the application generator were called into operation, with the result that the database software finally corrupted a table. This example thus demonstrates the action research equivalent to a Type II error, imputing support for a null hypothesis from the failure of experience to support the alternative hypothesis.

Since the time of their study's being published, the researchers have felt confidence in being able to apply their lesson beyond the particular action-research case is which it arose. They have formulated a general case in which prototyping management teams are inclined to reduce the probability rank of a known risk factor while the conditions that lead to such a risk became less favorable. Teams, in general, are theorized to rely on their own experience with risky aspects of projects, revealing a "guardian angel" mentality that inclines teams to ignore severely rising risk. The generality and confidence associated with this revised theory will remain ungauged until tested, like any other scientific theory.

Using the Schutz-inspired terminology of $\{\boldsymbol{T}, \boldsymbol{C'}, \boldsymbol{S_P'}, \boldsymbol{S_O'}\}$, we can say that the theory $\boldsymbol{T}$, at the beginning of the action research study, included the proposition, "The technical environment is unreliable." This led to the prediction $\boldsymbol{S_P'}$, that the phenomenon of corrupted database tables would occur. However, the observed state of affairs, $\boldsymbol{S_O'}$, was no corrupted database tables across the circumstances $\boldsymbol{C'}$ encountered in early cycles of the prototyping experience. Hence, because $\boldsymbol{S_P'}$ and $\boldsymbol{S_O'}$ differed, the academic members of the action research team judged the theory $\boldsymbol{T}$ to be falsified, and consequently revised it as $\boldsymbol{T_r}$ so as to include the revised proposition "The technical environment is reliable." The new form of the prediction, $\boldsymbol{S_P''}$, was that the phenomenon of corrupted database tables would not occur. However, further cycles of the prototyping experience encountered new circumstances $\boldsymbol{C''}$ where the functionality of the prototype software had become exceedingly complex. Here, in the latter stages of the prototyping experience, the observed state of affairs, $\boldsymbol{S_O''}$, included the occurrence of corrupted database tables; this contradicted the prediction $\boldsymbol{S_P''}$ and hence falsified the latest theory $\boldsymbol{T_r}$. As a result, the academic members of the action research team revised the theory once more so that, as $\boldsymbol{T_{rr}}$, it included the proposition, "The reliability of the technical environment enters as a moderating factor whose effect emerges upon the prototyping software's reaching a high level of complexity."

Note that, through these two cycles of deductive testing, the resulting theory T_{rr} can explain and be predictive across both sets of circumstances, C' and C'', hence achieving a greater degree of generality than possessed by either of its predecessor theories, T or T_r alone.

For another example of scientific deductive generalizing that similarly involves enhancing a theory's generality by revising it to be predictive across successive sets of circumstances, we turn to the 1998 study by Paré and Elam, who state: "the ultimate intent of this research is to build a new theory of IT [information technology] implementation" (p, 5). They applied the case research methods of Eisenhardt (1989) and Yin (1989) in observing IT implementation at Jackson Memorial Hospital, a large tertiary care teaching hospital in Miami, Florida. They conducted three case studies, the first one involving the additional of a new Electronic Medical Record system to the existing Patient Care System. Based in part of their observations in the first case, Paré and Elam formulated seven theoretical propositions about IT implementation, where the seventh proposition stated (p. 12):

P7: ***The implementation of an information system is characterized by a certain indeterminancy first reflected through the occurrence of unexpected challenges.***

They further explain (p.12): "This indeterminancy means that unexpected challenges, such as the data communication problems, are likely to be encountered. If no effective actions are taken to circumvent such challenges, unanticipated and likely undesirable consequences might result from the implementation process (e.g., unsolved data communication problems, disruption of MR [Medical Records] operations."

Paré and Elam then conducted a second case study, where they "systematically test the initial propositions developed in case one again evidence from case two" (p. 16). Regarding the seventh proposition, they note (pp. 17-18): "Finally, case two also confirms the contention that the implementation of an information system is likely to be characterized by a certain indeterminancy (P7). Contrarily to case one, though, such indeterminancy was not reflected through the occurrence of unexpected challenges, but rather was evidenced through a series of unanticipated and undesirable new consequences in the ICU (mainly caused by a laissez-faire implementation strategy)." To extend the generality of the seventh proposition so that it would also be able to apply to and predict these circumstances, Paré and Elam reformulated it as follows (p. 18, where underlining indicates a change to the proposition compared to its original form):

> **P7**: *The implementation of an information system is characterized by a certain indeterminancy first reflected through the occurrence of unexpected challenges. <u>This indeterminancy also means that the implemented system might not have all the effects originally envisioned by key actors</u>.*

Paré and Elam then conducted a third case study. This last part "systematically tests the confirmed and revised theoretical propositions against evidence from case three" (p. 21). Regarding the seventh proposition, they say (p. 23): "Lastly, case three confirms the contention that the implementation of an information system is likely to b characterized by a certain indeterminancy (P7). Similarly to what was found in case one, it appears that indeterminancy in IT implementation efforts can be reflected through the occurrence of unexpected challenges… Contrary to case one, though, the unexpected challenge encountered in case three was not the result of poor implementation practices. Rather, the search for a more advanced system was the result of an unexpected and uncontrollable event, namely, the vendor company's decision to discontinue the product." To extend the generality of the seventh proposition so that it would also be able to apply to and predict these latest circumstances, Paré and Elam reformulated it as follows (p. 24, where underlining indicates a change to the proposition compared to its previous form):

> **P7**: *The implementation of an information system is characterized by a certain indeterminancy first reflected through the occurrence of unexpected challenges <u>caused by either uncontrollable events or poor implementation strategies</u>. This indeterminancy also means that the implemented system might not have all the effects originally envisioned by key actors.*

We refer to the empirical circumstances that Paré and Elam observed in their first case study as $\boldsymbol{C'}$, those in their second case study as $\boldsymbol{C''}$, and those in their third case study $\boldsymbol{C'''}$, and we refer to Paré and Elam's original theory of IT implementation (including the original form of the seventh proposition) as $\boldsymbol{T}$, the revised theory (including the first revision of the seventh proposition, resulting from the second case study) as $\boldsymbol{T_r}$, and the final theory (including the last revision of the seventh proposition, resulting from the third case study) as $\boldsymbol{T_{rr}}$. <u>Note that, through these two cycles of deductive testing (the second and third case studies, where the first case study did not test, but formulated the theory in the first place), the resulting theory $\boldsymbol{T_{rr}}$ can</u>

explain and be predictive across the empirical circumstances of all three cases, C', C'', and C'', hence achieving a greater degree of generality than possessed by either of its predecessor theories, T or T_r alone. Hence, Paré and Elam achieve generality in their theory of IT implementation in a way entirely analogous to how Baskerville and Stage achieve generality in their theory of prototype development.

3. DISCUSSION

What is an appropriate way in which IS researchers may lay claim to generality for their research?

First, and most important, IS researchers should not give up claims to generality on the basis that their research involves a small "n" (for instance, a small number of organizations observed). The reason is that this would presume and reinforce the fallacy of the logic of induction – a logic rejected even by the natural-science model of research. A natural scientist who draws conclusions from an experiment does not rush to disclaim generality of her results on the basis that they are based on only a "single" experiment. In the same way, an IS researcher who draws conclusions from a single field site need not rush to disclaim the generality of his results.

Second, the analogy to the experimental natural scientist also reveals how an IS researcher may indeed lay claim to generality in her study. A natural scientist performing an experiment with favorable results (i.e., the theory's predictions are corroborated, not refuted) would need to state the details of the particular empirical circumstances for which the observed experimental results occurred. A claim of generality for the theory would mean that the theory can be expected also to hold in other instances *that share the same empirical circumstances*.[7] We note that an IS researcher who conducts a case study (or other intensive research) would be as able to do this as a natural scientist who conducts an experiment. And just as the natural scientist conducting an experiment would not need to apologize for her theory's not necessarily applying to empirical circumstances different from those in her experiment, the IS researcher who conducts a case study (or other intensive research) would not need to apologize for his

[7] Using the Schutz-based terminology, suppose we have $\{T, C', S_P', S_O'\}$, as we defined earlier. In the event that S_O' matches S_P', we can say that the theory T is sufficiently general for us to expect it to apply in all future occasions that instantiate the set of circumstances C'.

theory's not necessarily applying to empirical circumstances different from those in his field site. Along the same lines, future efforts to enhance the generality of the IS researcher's theory would not involve indiscriminately and randomly collecting more observations, but rather, would involve the targeting of additional specific empirical conditions to which the applicability of the theory is being questioned and, pending the observational results, would or would not be extended. In other words, a researcher does not enhance generalizability by simply increasing the number of supportive observations, but rather, by varying the range of different empirical circumstances in which the additional observations, consistent with the theory's predictions, are made.

Returning to some of the published examples of empirical research in the IS field that we presented in section 2.3, we suggest the following. We quoted Albers, Agarwal, and Tanniru (1994, p. 94) as saying: "From the point of view of theory development, while case studies provide useful anecdotal information, the generalizability from one specific instance to another is often limited." Our suggestion is that Albers, Agarwal, and Tanniru can say instead: "From the point of view of theory development, a case study should not be characterized as providing anecdotal information, but instead should be valued, first, for concretely demonstrating a specific instantiation of the circumstances in which the developed theory is known to apply and, second, for allowing additional applications of the same theory to other situations also involving instantiations of the same circumstances. Specifically, these circumstances are…" Similarly, we quoted Robey and Sahay (1996, p. 108) as saying: "Because they are drawn from a study of two organizations, these results should not be generalized to other contexts. Each context is different, so we should expect different contextual elements to interact with technical initiatives to produce different consequences." Our suggestion is that Robey and Sahay can say instead: "Because they are drawn from a study of two organizations, these results can apply to other contexts sharing the same circumstances of these two organizations. Where the elements in other contexts are the same in these two organizations, we can expect interactions of these elements with technical initiatives to produce the same consequences. Specifically, these circumstances are…" Finally, we quoted Hidding (1998, p. 311) as saying: "In particular, in-depth analysis of extensive data from only one organization reduces generalizability…" Our suggestion is that Hidding can say instead: "In particular, in-depth analysis from one organization enriches our knowledge of the fine details of the empirical circumstances in which the theory applies, hence increasing the reliability with which we can apply this theory in new settings in future occasions. Specifically, we can report the fine details of the empirical circumstances as comprising…"

Third, this notion of how to lay claim to generality also has an obvious ramification for practice. When an IS researcher demonstrates that her theory applies in certain circumstances, a practitioner would know that these are the circumstances in which he may trust the theory to work.[8] Where IS researchers simply abandon claims of generality and hence forego specifying what these circumstances are, the transfer of scientific findings to professional practice would be aborted.

4. CONCLUSION

The existing concept of "generalizability" is fastened on one peculiar form of generalizing. The concept is further confounded in the IS literature by its conflation with distinct concepts like generality and generalizing. Untangling this confusion reveals a variety of obscured, and sometimes legitimate forms, of generalizing. Many intensive research ventures could satisfy one of these forms, yet the authors self-flagellate over their inability to found their generalities on the one peculiar form that IS researchers have mindlessly idealized.

The self-flagellation is not merely annoying; it is harmful. When researchers unnecessarily divest their right to claim generality, their research audience is denied their analysis of the utility of their theories. When researchers generalize and claim generality, they encompass the larger scope of phenomena, beyond those directly captured by their research, to which their findings and understandings apply (Babbie, 1990). By renouncing their right to generalize and claim generality, intensive researchers lose the latitude to explain the wide field of uses for their findings.

The IS research field is relatively vocational and operates in concert with technologies that are incredibly fast-moving. Our understanding of the social and organizational aspects of our field may be trailing far behind our grasp of the technical issues. Unnecessarily confining the application of new theories from intensive research is helping to cripple our ability to keep pace.

[8] Using the Schutz-based terminology, the same circumstances would be the $\boldsymbol{C'}$ in $\{\boldsymbol{T}, \boldsymbol{C'}, \boldsymbol{S_P'}, \boldsymbol{S_O'}\}$ for the information systems researcher engaging in scientific generalizing, and the $\boldsymbol{C'}$ in $\{\boldsymbol{A'}, \boldsymbol{C'}, \boldsymbol{S'}\}$ for the practitioner engaging in everyday generalizing.

Applying a variety of forms of generalizing in IS promises to improve the development of more current and more useful social and organizational theory.

REFERENCES

Albers, M., Agarwal, R., and Tanniru, M. (1994) "The Practice of Business Process Reengineering: Radical Planning and Incremental Implementation in an IS Organization" in *SIGCPR '94. Proceedings of the 1994 Computer Personnel Research Conference on Reinventing IS: Managing Information Technology in Changing Organizations*, ACM Press, Arlington, VA, 87-96.

Argyris, C., and Schön, D. (1978) *Organizational Learning: A Theory of Action Perspective*, Addison-Wesley, Reading.

Argyris, C., and Schön, D. (1991) "Participatory Action Research and Action Science Compared" in *Participatory Action Research*, (W. F. Whyte, ed.), Sage, Newbury Park, NJ, 85-96.

Babbie, E. (1990) *Survey Research Methods 2nd Ed.*, Wadsworth, Belmont, California.

Baskerville, R. L., and Stage, J. (1991) "Developing the Prototype Approach in Rapid Systems Modelling." *Institute for Electronic Systems R91-35*, September 1991, The University of Aalborg, Aalbor, Denmark.

Baskerville, R., and Stage, J. (1996) "Controlling Prototype Development Through Risk Analysis." *MIS Quarterly*, 20 (4), 481-504.

Berger, P., and Luckmann, T. (1966) *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*, Anchor Press, New York.

Brown, C. (1997) "Examining the Emergence of Hybrid IS Governance Solutions: Evidence from a Single Case Site." *Information Systems Research*, 8 (1), 69-94.

Copi, I. (1986) *Introduction to Logic*, Macmillan, New York.

Eisenhardt, K. (1989) "Building Theories from Case Study Research." *Academy of Management Review*, 4, 532-550.

Hidding, G. J. (1998) "Adoption of IS Development Methods across Cultural Boundaries," in *Proceedings of the Nineteenth International Conference on Information Systems*, (R. Hirschheim, M. Newman, and J. I. DeGross, eds.), 308-312.

Hult, M., and Lennung, S. (1980) "Towards a Definition of Action Research: A Note and Bibliography." *Journal of Management Studies*, 17, 241-250.

James, W. (1975) "What Pragmatism Means" in *Pragmatism*, (W. James, ed.), Harvard University Press, Cambridge, MA, 27-44.

Jarvenpaa, S. L., and Leidner, D. E. (1997) "An Information Company in Mexico: Extending the Resource-Based View of the Firm" in *Proceedings of the Eighteenth International Conference on Information Systems*, (K. Kumar and J. I. DeGross, eds.), 75-87.

Lee, A. S. (1989) "A Scientific Methodology for MIS Case Studies." *MIS Quarterly*, 13 (1), 33-50.

Lee, A. S. (1999) "Researching MIS," in *Rethinking MIS* (W. L. Currie and R. D. Gallier, eds.), Oxford University Press.

Nandhakumar, J., and Jones, M. (1997) "Designing in the Dark: the Changing User-Developer Relationship in Information Systems Development" in *Proceedings of the Eighteenth International Conference on Information Systems* (K. Kumar and J. I. DeGross, eds.), 75-87.

Orlikowski, W. (1993) "Case tools as organizational change: investigating incremental and radical changes in systems development." *MIS Quarterly* 17 (3), 309-340.

Oxley, D., and Curtis, B. (1997) "Reliable Work on a Death March Schedule." *American Programmer*, 10 (2), 12-15.

Paré, G. and Elam, J. (1998) "Broadening Our Understanding of IT Implementation: Toward the Development of a Process Theory." Cahier GreSI no. 98-01, Groupe de recherche en systèmes d'information, École des Hautes Études Commerciales, Montréal, Québec.

Popper, K.R. (1968) *The Logic of Scientific Discovery*, Harper Torchbooks, New York.

Roberts, S. M. (1997) "Surviving and Succeeding in a Death March Project." *American Programmer*, 10 (2), 12-15.

Robey, D., and Sahay, S. (1996) "Transforming Work through Information Technology: A Comparative Case Study of Geographic Information Systems in County Government." *Information Systems Research*, 7 (1), 93-110.

Schultze, U. (forthcoming) "A Confessional Account of an Ethnography about Knowledge Work." *MIS Quarterly*.

Schutz, A. (1973) "Common Sense and Scientific Interpretation of Human Action" in *Alfred Schutz, Collected Papers I: the Problem of Social Reality*, Martinus Nijhoff, The Hague.

Stein, E. W., and Vandenbosch, B. (1996) "Organizational Learning During Advanced System Development: Opportunities and Obstacles." *Journal of Management Information Systems*, 13 (2), 115-.

Strauss, A., and Corbin, J. (1990) *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*, Sage Publications, Newbury Park, CA.

Yin, R. (1989) *Case Study Research, Design and Methods*, Sage Publications, Beverly Hills, CA.

Yourdon, E. (1997) "Editorial." *American Programmer*, 10 (2), i-1.