

December 2006

Other-Settings Generalizability in IS Research

Peter Seddon
University of Melbourne

Rens Scheepers
University of Melbourne

Follow this and additional works at: <http://aisel.aisnet.org/icis2006>

Recommended Citation

Seddon, Peter and Scheepers, Rens, "Other-Settings Generalizability in IS Research" (2006). *ICIS 2006 Proceedings*. 70.
<http://aisel.aisnet.org/icis2006/70>

This material is brought to you by the International Conference on Information Systems (ICIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ICIS 2006 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

OTHER-SETTINGS GENERALIZATION IN IS RESEARCH

Quantitative Research

Peter B. Seddon

University of Melbourne
p.seddon@unimelb.edu.au

Rens Scheepers

University of Melbourne
r.scheepers@unimelb.edu.au

Abstract

This paper presents a simple conceptualization of generalization, called other-settings generalization, that is valid for any IS researcher who claims that his or her results have applicability beyond the sample where data were collected. An other-settings generalization is the researcher's act of arguing, based on the representativeness of the sample, that there is a reasonable expectation that a knowledge claim already believed to be true in one or more settings is also true in other clearly defined settings. Features associated with this conceptualization of generalization include (a) recognition that all human knowledge is bounded, (b) recognition that all knowledge claims—including generalizations—are subject to revision, (c) an ontological assumption that objective reality exists, (d) a scientific-realist definition of truth, and (e) identification of the following three essential characteristics of sound other-settings generalizations: (1) the researcher must clearly define the larger set of things to which the generalization applies; (2) the justification for making other-settings generalizations ultimately depends on the representativeness of the sample, not statistical inference; (3) representativeness is judged by comparing key characteristics of the proposition being generalized in the sample and target population. The paper concludes with the recommendation that future empirical IS research should include an explicit discussion of the other-settings generalizability of research findings.

Keywords: Research methodology, other-settings generalization, OSG, external validity

Introduction

This paper presents a simple conceptualization of generalization, called other-settings generalization, that is valid for any Information Systems (IS) researcher who claims that his or her results have applicability beyond the sample where data were collected. The reasons that a paper on this topic is needed are as follows. First, as Lee and Baskerville (2003) point out, IS is an applied discipline. It is therefore important that the likely implications for practitioners of IS research findings be laid out as clearly as possible. Second, with increasing globalization of IS research, it is likely that due to important cultural and institutional differences, findings from studies in one country or culture may not apply directly in other countries. The applicability of findings between different contexts therefore needs to be discussed clearly. Third, all knowledge claims are bounded. Therefore, as Dubin (1969) and Whetten (1989) point out, rigorous research must identify the boundaries of applicability of any theories. Fourth, there is a mounting body of evidence, e.g. see Lee and Baskerville (2003) and King and He (2005), that many published IS research reports, even in the top journals, handle generalizations of knowledge claims poorly.

With regard to the fourth point above, many quantitative researchers collect survey data with response rates of the order of 30% (King and He 2005, Table 5), yet proceed to report p-values from statistical packages as if they are automatically meaningful. Equally, as Lee and Baskerville (2003) have shown, many case-study researchers issue disclaimers that their findings are not generalizable to other settings, although many such findings clearly do have relevance to other settings. Turning to the review process, our experience has been that many reviewers draw unsound conclusions about generalizability of research findings, and reject papers based on invalid rules of thumb such as “you can't generalize from a single case”, or “you can't generalize from a sample of students”, or “I am not convinced that the non-random, non-representative nature of the evidence is going to result in valid findings”. In short, the handling of generalizability issues in IS research is challenged in many dimensions.

To address these issues, and to add weight to the arguments of Lee and Baskerville (2003) and King and He (2005), this paper presents a summary of the fundamentals of generalization that integrates, simplifies, and in some cases re-organizes the work of many prior researchers. None of the individual points we raise in this paper is fundamentally new, but the aggregation *is* new. We call this aggregation of ideas, defined in Table 1, “other-settings generalization” (OSG). “Other settings” include both the population from which the sample was or might conceptually have been drawn, and other populations beyond the sample population, e.g. those in other countries or cultures where the results of the study might also apply. OSG is a very important subset of the broad 2x2 framework presented by Lee and Baskerville (2003)¹. This subset is important because when most researchers talk of the generalizability or external validity of research findings, e.g. King and He (2005), the concept they seem to have in mind is other-settings generalizability.

Table 1. Summary of Key Characteristics of Other-Settings Generalization (OSG)
<p><i>Definition</i></p> <p>An other-settings generalization is the researcher’s act of arguing, based on the representativeness of the sample, that there is a reasonable expectation that a knowledge claim already believed to be true in one or more settings is also true in other clearly defined settings.</p>
<p><i>Key Characteristics of sound other-settings generalizations</i></p> <ol style="list-style-type: none"> 1. All human knowledge is bounded 2. All knowledge claims, including generalizations, are subject to revision 3. Objective reality exists (this is a soft-positivist ontological stance) 4. Truth beyond reasonable doubt is sufficient (this is a scientific-realist definition of truth) 5. The only sound basis for generalization of knowledge claims is the representativeness of the sample(s) on which the knowledge claim is based 6. Representativeness is demonstrated by comparing key characteristics of the proposition being tested in the sample and target population

It is important to draw some boundaries for this paper. First, although other-settings generalizations are of interest in all the social sciences, in this paper we restrict our attention to other-settings generalizations in information systems. We do this because the publications we examine in the next section come only from the IS literature. We do, however, revisit the question of the generalizability of our claims to other social-science disciplines, towards the end of this paper. Second, we also assume that readers are familiar with the principles of inferential statistics, including the concepts of sample, sampling frame, and population, and understand that p-values (e.g. “**”, $p < 0.05$) calculated by most statistical packages assume data have come from a random sample and are *not* automatically meaningful if the sample studied is *not* a random sample.

Returning to our main argument, according to Weber (2004, p. vi), the objective of all researchers is to “enhance their understanding of the world (whatever the world might be)”. However, attempts to understand how the world works are inevitably incomplete because researchers have limited cognitive capabilities and only localized understandings of what they and prior researchers see and have seen. Recognizing these limitations, Dubin (1969) and Whetten (1989) suggest that researchers who build theories and models to explain how the world works should specify clearly the boundaries to their knowledge claims. Surprisingly, however, few IS researchers attempt to delineate clearly the boundaries beyond which their research findings might not apply. To improve this situation, we argue that IS researchers, both positivists and interpretivists, should normally be *expected* to include a discussion of the potential generalizability, and limits to generalizability, of findings of any empirical study in any research paper.

The remainder of this paper is structured in three parts. First, we summarize the mounting evidence that generalization is poorly handled in IS research. Second, in the main body of the paper, we present the aggregation

¹ Lee and Baskerville’s (2003) framework accommodates many different types of concepts that various authors have called generalization. Our conceptualization of OSG corresponds to Lee and Baskerville’s second types of EE and ET generalization “beyond the sample or domain from which the researcher has actually collected data” (p.233).

of ideas that we would like to associate with the label “other-settings generalization” (OSG). These are the ideas summarized in Table 1. Finally, in a relatively brief section, we present suggestions for improving discussion of OSG in future IS research.

Problems with Generalization in IS Research

Two recent papers—Lee and Baskerville (2003) and King and He (2005)—provide an indication of some unsound practices in the treatment of generalization in leading IS research journals and conferences. With respect to case studies, Lee and Baskerville (2003) identify a dozen studies where case-study researchers understate the likely generalizability of their results. With respect to surveys, King and He (2005) show that many survey-based researchers mishandle generalizability claims by ignoring the fact that samples they have analyzed are unlikely to be representative of the populations they (often implicitly) claim to have studied. For both types of research, we argue, a discussion of other-settings generalizability would provide the reader with access to useful insights the researchers may have about the likely limits in the applicability of their results to settings other than those on which the findings are based.

Case-study Research Under-claiming Generalizability

In their paper on generalizing in IS research, Lee and Baskerville (2003) present a list of 12 case-study research papers (Table 1, p.223) and state that in each case the researchers involved have applied a “statistical sampling-based conception of generalizability” “inappropriately” in discussing their research. What they mean by “inappropriately” is that the researchers have under-claimed the general applicability of their conclusions. The reason for thinking that reporting in these 12 case-study-based IS research papers should be improved is that despite their disclaimers about the lack of generalizability from their case studies, all twelve groups of researchers (some of whom were using interpretive research methods) wrote as if they believed their findings *were* relevant to other settings. In some cases, they made much stronger claims to generality in their papers than their disclaimers implied.

For example, Robey and Sahay’s (1996) paper on the consequences of GIS adoption in their case studies of two US counties contains a typical disclaimer:

“Because they are drawn from a study of two organizations, these results should not be generalized to other contexts.” (p.108)

However, their language in discussing their findings speaks otherwise:

“Because the same technology (Arc/Info) was experienced differently during its introduction in the two counties, the results reported here strongly support the idea that information technology’s consequences are socially constructed, i.e., that technology’s social consequences depend on its social meanings more than on its material properties.” (p.106)

The implication of the latter statement, and, in fact, the strongest argument running through Robey and Sahay’s entire paper, is that “information technology’s consequences are socially constructed”. Furthermore, the tone of the paper suggests that this is true for *all* applications of IT in *all* organizations (not just the two US county offices that were using ESRI’s ARC/INFO GIS software). If Robey and Sahey had made this claim explicitly it would have been a sweeping other-settings generalization (moreover, one that we believe could be justified beyond reasonable doubt)! There are similar, if less strong, mixed messages in Lee and Baskerville’s other 11 studies. In fact, all 12 papers make quite strong general claims, whilst simultaneously issuing caveats that their findings could not be generalized from the small number of cases. We believe that in many cases the stronger claims are justified.

External-validity Problems in Survey-based IS Research

Campbell and Stanley’s (1963) term “external validity” is frequently used to describe the validity of generalizations from samples to populations. In their study of external validity in survey-based research, King and He (2005) document what they call “coverage error” and “nonresponse error” in 199 empirical papers published in three top journals, *MIS Quarterly* (MISQ), *Information Systems Research* (ISR), and the *Journal of Management Information Systems* (JMIS) over the five years from April 1999 to April 2004. According to King and He, coverage error is “introduced when the frame from which the sample is drawn does not include all relevant characteristics of the

population to which inferences are to be drawn” (King and He 2005, p.882), and non-response error is “introduced if non-respondents are different from respondents in terms of characteristics that are relevant to the study [Dillman 2000]” (King and He 2005, p.885).

King and He’s Table 3 shows that “over 40% of papers in all three journals ignore coverage error and almost 20% only mention it briefly” (p.887). Further, their Table 6 shows that for 81 mail surveys, “well over half of all studies in all three journals (approx. 48%-70%) neither mentioned nor performed nonresponse analysis” (pp.888-9). These are large percentages. Since both coverage and nonresponse error are sources of problems in generalizing, King and He’s findings provide clear evidence that the state of IS generalization practice in survey-based research is not as satisfactory as it should be. However, King and He do not discuss some key ideas that are at least as important for making sound generalizations in IS research. These include (a) the idea that all knowledge claims are bounded, and (b) the fact that p-values calculated by statistical packages are not meaningful unless samples of data studied are probability samples (which most survey studies are not). In addition, they do not consider case studies. Therefore, we decided to conduct our own analysis of published papers in two top IS journals.

Other-settings Generalizability in Recent IS Research

In our analysis, we examined all publications in MISQ and ISR in the two years, 2003 and 2004, to assess how they treated the generalizability of their knowledge claims². During these two years, there were 38 papers published in ISR and 45 in MISQ, a total of 83 papers. As King and He (2005) also report, clues about generalizability claims were usually in: (a) the Methodology section, where authors explained why they had chosen the sample they had, and, in the case of surveys, why non-response bias was (usually) not a problem in their study, and/or (b) the Discussion, Limitations³, Implications for researchers, Implications for management, and/or Conclusion sections, where authors frequently discussed why their results might be generalizable or not, or offered advice to practitioners on how results from their study might be useful. Often, e.g. in Kohli and Kettinger’s (2004) action-research study or Bassellier et al.’s (2003) survey, it was just assumed that the relevant other setting was something like “all users or managers in North American companies” and that readers understood this to be the case.

Results from our analysis are shown in Table 2. We judged that 72 of the 83 papers were interested in some way with the generalizability of their findings to other settings. Analytic studies and Simulations, e.g. Fan et al. (2003), van der Aalst and Kumar (2003), and Raghu et al. (2004) were included in the 72 because these papers claim to model phenomena in the real world. Reviews, such as Fichman (2003) and Melville et al. (2004), were also included in the 72 because in integrating the literature, the authors are building theories and propositions that they believe apply to practice. Overall, in the 71% of the papers where other-settings generalizability was judged relevant, the authors wrote as if their results did have generalizability to other settings, yet only 46% of the papers discussed generalizability at all, and only 17% made any attempt to discuss the boundary conditions beyond which their theories or findings might not apply.

Summary: The State of Generalization Practice in Empirical IS Research

Taken together, the above three studies are consistent in showing that generalization is poorly handled in IS Research. Case-study researchers who claim that their results cannot be generalized are under-claiming the relevance of their work. Quantitative researchers who analyze samples and do not discuss boundary conditions of their knowledge claims are implicitly over-claiming generalizability of their research. Further, in our experience, reviewers who reject papers based on implicit judgments about usefulness of samples for drawing “general” conclusions also often base those conclusions on unsound logic.

² We deliberately restricted the sample to just two years to make the point that if the sample is carefully chosen, valid general claims can be based on relatively small samples. The question is: Would the percentages in Table 1 be expected to change much if the sample were for, say, five years, or for a broader set of top-tier IS journals?

³ In our sample of 83 studies, the discussion of limitations section was usually a mix of internal and external validity (generalizability) issues. Despite the importance of the latter, discussion of internal validity issues often dominated.

Table 2. Discussion of Generalizability in Papers Published in *ISR* and *MISQ* in 2003 and 2004

Method	Total	<i>ISR</i>	<i>MISQ</i>	Other-settings generalizability assumed or implied	Generalizability discussed	Boundary conditions for knowledge claims discussed
<i>Papers where generalizability to other settings is relevant</i>						
Survey	18	7	11	11	11	5
Experiment	13	8	5	8	10	1
Analytic	10	10	0	6	0	0
Case study	8	3	5	7	6	3
Review	6	1	5	6	2	1
Action Research	5	0	5	3	0	0
Archival data analysis	5	2	3	4	2	2
Simulation	3	3	0	3	0	0
Meta analysis	2	1	1	2	2	0
Opinion	2	0	2	1	0	0
Total	72	35	37	51	33	12
	100%			71%	46%	17%
<i>Papers where generalizability to other settings is not relevant</i>						
Research method	7	3	4	not applicable		
What is IS?	2	0	2	not applicable		
Design idea	1	0	1	not applicable		
Ethics	1	0	1	not applicable		
	11	3	8			
	83	38	45			

What is needed to solve these problems are, first, a clear understanding of what generalization means and why generalization matters in IS research, and second, some clear recommendations on the way generalization issues should be discussed in future IS research. In the remainder of this paper, we address both these needs. As noted earlier, our conceptualization of generalization is a subset of the all-encompassing 2 x 2 framework presented by Lee and Baskerville (2003). However, we believe that this is the key subset to understand, because when most researchers talk of the generalizability or external validity of research findings, the concept they have in mind is other-settings generalizability.

Other-Settings Generalization

The world is interested in generalizations that are true. In this section we define “truth” and “other-settings generalization”. But before defining these terms, three tightly coupled questions must be answered:

1. What is the ontological stance of this paper?
2. What things do social science researchers attempt to generalize to other settings?
3. Who makes other-settings generalizations, the writer or the reader of a research report?

First, the ontological stance of this paper is depicted in Figure 1. Like Kirsch (2004), who describes her ontological position as “soft positivism”, our position is somewhere in the middle band between extreme interpretivism and extreme positivism⁴. Specifically, we believe that objective reality exists beyond the human mind, though our perceptions about that reality are inextricably bound to the stream of experiences we have had throughout our lives. Further, we believe that there are many regularities and patterns in this objective reality that researchers seek to uncover, but that these regularities and patterns tend to apply in only limited contexts and are likely to be different for different types of people (managers, teenagers, etc.), different cultures⁵, and over time. These beliefs underpin the remainder of the arguments in this paper.

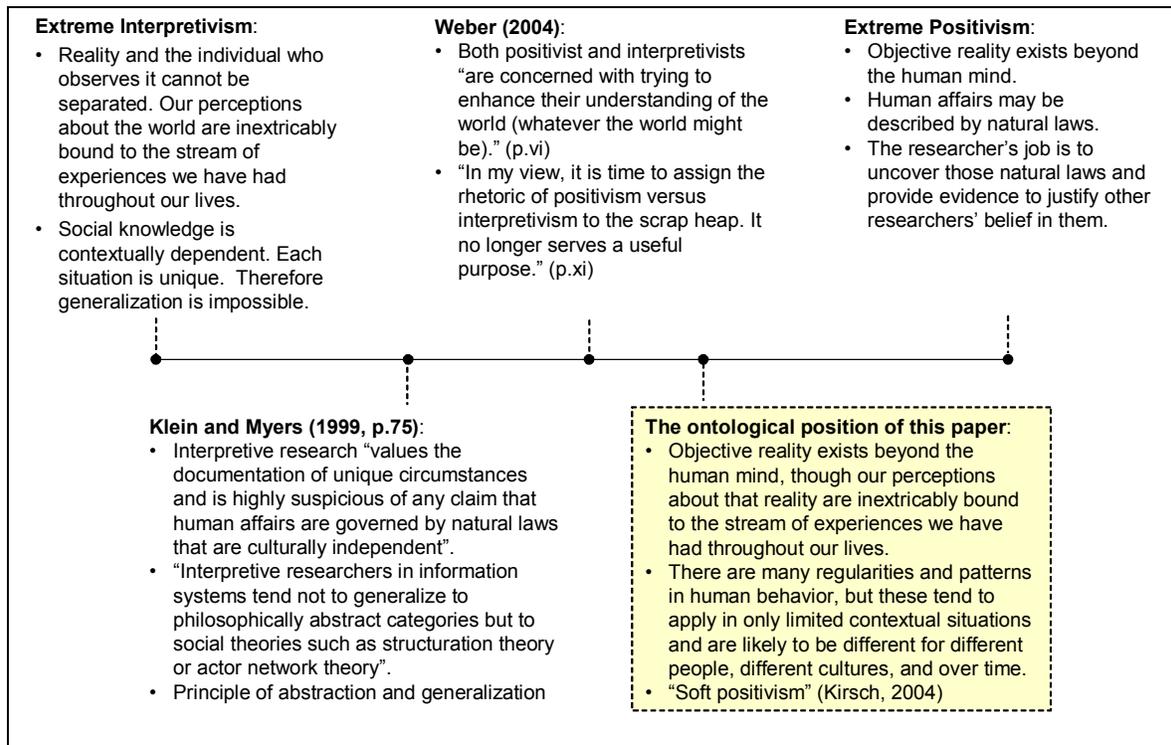


Figure 1. Reconciling the Extreme Ontological Positions of Interpretivist and Positivist Research on the Generalizability of Research Results

Second, we believe that the things that social-science researchers seek to generalize to other settings are *knowledge claims* from their research studies. Such knowledge claims are usually presented in the form of propositions, models or theories⁶ that the researchers believe explain the facts in both past research and their present studies. There is always uncertainty about the truth of such sample-based findings. Further, in the course of normal science, it is expected that over time, most theories will be replaced by new, better theories. For example, UTAUT (Venkatesh et

⁴ Hunt (2003) devotes twenty pages (pp.198-218) to discussing what he describes as misconceptions about differences between positivism and interpretive/qualitative research. He comes to a conclusion similar to Weber (2004), shown in the middle of Figure 1.

⁵ For instance, due to different cultural mores (Leidner and Kayworth 2006), a theory that explains management behavior in corporate America today might not be valid for explaining management behavior in China or Thailand.

⁶ Many definitions of theory exist (Gregor 2006). According to Neuman (2003: 50): "Theories contain concepts, their definitions, and assumptions. More significantly, theories specify how concepts relate to one another. Theories tell us whether concepts are related, and if they are, how they relate to each other. In addition, theories state why the relationship does or does not exist."

al. 2003) seems likely to replace TAM (Davis 1989) as our current best explanation of individual-user IT adoption and use, at least for the settings in which it and TAM have been tested so far.

Note that our position in this paper on “what is generalized” differs from that of Baskerville (1996), who describes generalization as a two-stage process where (1) findings from a sample of things (which Baskerville terms the “base case”) are used to create a “general case”, then (2) characteristics of the general case are applied to predict outcomes in some “goal case”. Under Baskerville’s view, which is a valid alternative to our conceptualization of generalization,:

“A general case is an abstract, theoretical case that manifests a relevant subset of the characteristics of the base case (or shared characteristics of the base cases).” (p.9)

In this paper, the things being generalized are sample-based knowledge claims, e.g. propositions, models, and theories. We do not conceptualize generalization in terms of some typical or base case.

Third, we follow Palmquist et al. (2004) in distinguishing between “generalization”, meaning decisions made by *researchers* about the generalizability of their own findings to other settings, and “transferability”, meaning decisions made by *readers* about the applicability of someone else’s research findings to their own organization. According to Palmquist et al. (2004):

“Transferability is a process performed by *readers* of research. Readers note the specifics of the research situation and compare them to the specifics of an environment or situation with which they are familiar. If there are enough similarities between the two situations, readers may be able to infer that the results of the research would be the same or similar in their own situation. In other words, they ‘transfer’ the results of a study to another context.”

Truth

As Hume (1739), Popper (1992), Dubin (1969), Shadish et al. (2002), Hunt (2002, 2003), Lee and Baskerville (2003), and many other philosophers and researchers point out, claims about generalizability from samples to entire populations can never be proven true in the true/false sense of formal logic:

“No matter how many instances of white swans we may have observed, this does not justify the conclusion that all swans are white” (Popper 1992, 27).

Faced with the logical impossibility of proving generalizations from samples to other settings true, if we as researchers wish to generalize—and developments in science over many centuries have shown repeatedly that attempts to make such generalizations have been useful—we have to soften our test of truth of an other-settings generalization.

We have no desire to enter into a philosophical debate on the meaning of truth, but since our definition of other-settings generalization relies on the term, we need to explain what we mean by “truth” in this paper. In this paper we adopt Hunt’s (2003, 2003) scientific-realist position on truth⁷, and define the truth of an other-settings generalization as follows:

⁷ Hunt (2003) devotes an entire 300-page book to a review of various philosophies of science—from the early Greek philosophers (Socrates, Aristotle) through what he terms the “classical empiricists” of the sixteenth to nineteenth centuries (Bacon, Hume, Mill), the “classical realists” (Russell, Wittgenstein), “pragmatists” (Pierce), “logical positivists” (Schlick, Carnap), “falsificationists” (Popper), “historical relativists” (Kuhn, Feyerabend), and “scientific realists” (Harre) of the twentieth century—and concludes by supporting Harre’s (1986) view that many philosophers have fallen into the trap of defining truth in such absolute terms that no knowledge claims can be described as true. In his Table 7.1, Hunt (2003, p.239) classifies 24 different philosophical positions on truth—e.g. from the deconstructive postmodern to the theocratic—and concludes in favor of scientific realism’s position on truth:

“In short, scientific realism proposes that (1) the world exists independently of its being perceived (classical realism), (2) the job of science is to develop genuine knowledge about that world, even though such knowledge will never be known with certainty (fallibilistic realism), and (3) all knowledge-claims must be critically

An other-settings generalization statement is true if it is True (in the binary, true/false sense of formal logic) beyond reasonable doubt, though such a statement is always to subject to later disconfirmation.

Using this definition of truth as our criterion for accepting other-settings generalizations allows us to avoid the logical consequences of what Lee and Baskerville (2003) call Hume's truism⁸. Since the propositions and theories we start with in generalizing are never True in the true/false sense of formal logic, and are likewise always subject to uncertainty and revision this is not an unreasonably limitation. Truth beyond reasonable doubt is as close as any researcher can get to formal logic's concept of Truth.

Other-settings Generalization

Although the English word "generalization" has many meanings, the generalization concept of interest in this paper is other-settings generalization:

An other-settings generalization is the researcher's act of arguing, based on the representativeness of the sample, that there is a reasonable expectation that a knowledge claim already believed to be true in one or more settings is also true in other clearly defined settings.

In addition to using the word "true", defined above, the preceding definition also uses the word "representativeness". Judgments about representativeness are aided by considering Shadish et al.'s (2002) five principles, discussed shortly. As a guide, the intuition of what we mean by "representative" is most easily conveyed using sampling logic (which, of course, does not necessarily apply to case studies) as follows:

For the variables relevant to the study, statistics calculated from a representative sample are similar to those of the population from which the sample was selected.

Other Authors' Concepts Corresponding to Other-settings Generalization

Table 3 shows some leading authors' terms for other-settings generalization, including Campbell and Stanley's (1963) frequently used term "external validity". We would happily use their term (external validity) instead of introducing a new one (other-settings generalization), but "external validity" describes the *degree* of truth of other-settings generalizations, and is not readily applied to the *process* of generalizing. Further, each term in Table 3 has so many ideas related to it that it seems wiser to coin a new term to describe what we argue is a particularly useful aggregation of ideas on generalization.

The only contentious claim in Table 3 is that, despite Yin's (1989, 1994, 2003) claims to the contrary, we argue that Yin's statistical and analytical generalization are just two alternative logical pathways for justifying other-settings generalization. Arguments supporting this contention are presented in Appendix 2.

evaluated and tested to determine the extent to which they do, or do not, truly represent or correspond to that world (critical realism).

⁸ Lee and Baskerville (2003: 224-228) present a four-page discussion of Hume's truism. The problem is well illustrated by the following example of inductive logic from Wood (2000):

1. In past experience, all Fs have been Gs.
2. Therefore, all future Fs will be Gs, or the next F will be a G.

Unless one makes the further assumption that nature is uniform, proposition 2 cannot validly be drawn from proposition 1. Further, any attempt to prove that nature is uniform suffers from the same problem: knowledge that nature has been uniform in the past does not ensure that it will be in future. Hence Hume's truism: "induction or generalization is *never* fully justified logically" (Campbell and Stanley, 1963: 17, emphasis added).

Table 3. Leading Authors' Terms for Other-settings Generalization

Author(s) (chronological order)	Concept corresponding to other-settings generalization
Campbell and Stanley (1963)	Campbell and Stanley's <i>external validity</i> describes the degree of truth of other-settings generalizations. This same external validity concept is discussed in later editions of the book, i.e., in Cook and Campbell (1979) and Shadish, Cook and Campbell (2002).
Dubin (1969)	Generalizing "to a larger domain".
Cronbach (1982)	Cronbach (1982) suggests that discussions of generalization should be organized around four concepts: units (<i>U</i>), treatments (<i>T</i>), observations (<i>O</i>), and settings (<i>S</i>). Cronbach's extrapolation to UTOS or *UTOS corresponds to our concept of other-settings generalization.
Yin (1989, 1994, 2003)	Yin's concepts of statistical and analytical generalization are two alternative mechanisms for justifying other-settings generalization. Justification for this claim is provided in Appendix 2.
Walsham (1995)	Walsham's third type of generalization, "drawing of specific implications" (p.79), corresponds to other-settings generalization. He suggests that the following statement from Walsham and Waema (1994, p.171) is a generalization from the single interpretive case study: "An ad hoc methodological approach to the development of computer-based information systems, accompanied by a clear business focus, can lead to rapid systems development, but the price paid for such an approach can be inflexibility and a lack of adequate integration". Note the word "can" not "will". The implication is that this statement is true in other (unspecified) settings.
Klein and Myers (1999)	Klein and Myers' "theoretical general concepts that describe the nature of human understanding and social action" in the second half of the Principle 4 appear to be other-settings generalizations. To illustrate their Principle 4, Klein and Myers discuss the use of Latour's actor-network theory in a paper by Monteiro and Hanseth (1996). They say: "According to this theory, humans and non-humans are linked together into actor networks. The theory assumes that actors pursue interests, and that these interests can become inscribed in technical or social arrangements." (Klein and Myers 1999, p.76). From an other-settings-generalization perspective, the preceding quotation describes a knowledge claim asserted to be true for <i>all human social action</i> .
Lee and Baskerville (2003)	Lee and Baskerville's (2003) paper on generalization is the most comprehensive summary to date of research thinking on generalization in the IS literature. Their paper is intended to encompass the "many different concepts in the scholarly discourse on generalizability" (p.241). Lee and Baskerville's second types of EE and ET generalization, which involve generalization "beyond the sample or domain from which the researcher has actually collected data" (p.233), are other-settings generalizations.

The Process of Other-settings Generalization

The process followed in other-settings generalization is depicted in Figure 2. Suppose that a proposition, model or theory, *x*, is believed to be true (i.e., beyond reasonable doubt) in one or more settings, e.g. in a case study in a single firm, as the result of a single experiment, or from an analysis of survey data (Figure 2, step 1). As a result of some stimulus (Figure 2, step 2) the observer formulates the view that there is a reasonable expectation that *x* is true in other clearly specified settings beyond the scope of the initial study (Figure 2, step 3). At that moment, the observer has generalized to other settings.

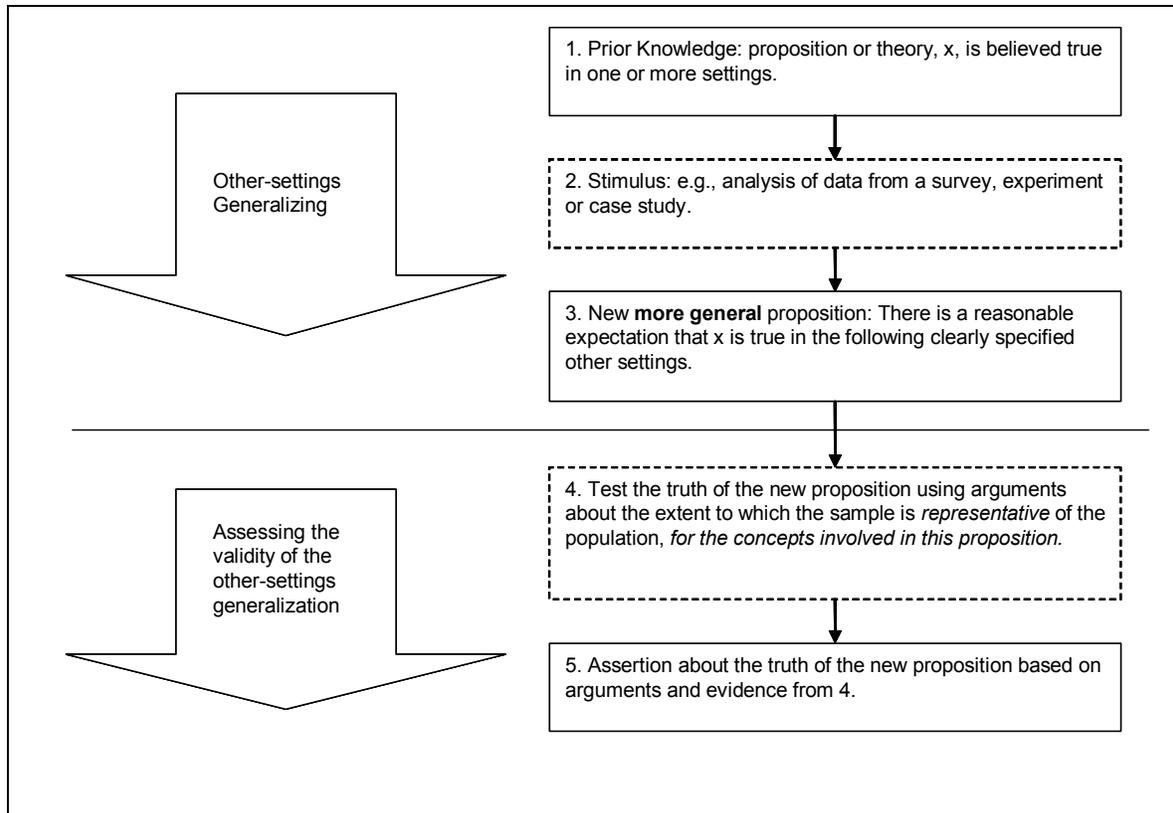


Figure 2. The Process of Other-settings Generalization

Under the view depicted in Figure 2, an “other settings” generalization is simply a shift in one’s mental view of the world: one asserts, based on the evidence to hand, that the phenomenon one believes to be true in one setting (e.g. in a case study in a single firm, multiple case studies, based on data from a survey, or from one or more experiments) is likely to be true in one or more other settings that one has not observed (e.g. other firms where similar contextual factors also exist, or the population of things from which the sample used in the experiment was drawn). However, not all such assertions are true beyond reasonable doubt; they need to be tested. This is where steps 4 and 5 in Figure 2 come in. The *only* techniques human beings have for assessing the truth of other-settings generalization assertions are logic and more empirical testing. One has to rely on induction, assumptions about the uniformity of nature, arguments about the “representativeness” of the sample, and subsequent empirical testing to provide support for one’s other-settings generalizations⁹. (Inferential statistics and p-values are helpful *only* if the sample is a probability sample, and in practical research—including surveys, experiments, and case studies—few samples are. Appendix 1 shows that inferential statistics and p-values may be useful if the sample is *representative* of things in those other settings.) Representativeness is therefore very important. Criteria for claiming representativeness are discussed in the next section.

⁹ In this respect Weick (1984, p.117), citing McGuire (1983), states that “...empirical confrontation is not a test of whether a theory is correct; rather, it is a discovery process: to make clear what the theory means, disclose its hidden assumptions, and clarify the conditions under which it is true or false. Through a series of studies, one seeks to establish the theory's pattern of adequacy, that is, to articulate the settings in which its inevitable misrepresentations are tolerable or intolerable.”

Shadish et al.'s Five Principles for Generalizing

The clearest guidelines that we are aware of for assessing representativeness come from Cronbach (1982) and Shadish et al. (2002). Here we summarize Shadish et al.'s five principles for generalizing, which build on Cronbach's Units, Treatment, Observations and Settings (UTOS) framework¹⁰.

Shadish et al. (2002) suggest "that scientists make causal generalizations in their work by using the five closely related principles" shown in Table 4. Each of the five principles in Table 4 is discussed in detail by Shadish (1995), and Shadish et al. (2002), so the reader can refer to them for more details. In essence, however, Shadish et al. (2002) argue that since no two things are exactly identical, judgments about whether one set of things is similar to, or representative of, some other set of things, must be based on establishing that they are similar with respect to attributes that matter.

Table 4. Shadish et al.'s (2002: Table 11.1, pp. 357-8) Five Principles for Generalizing

1. *Surface Similarity*. Assessing the apparent similarities between study operations and the prototypical characteristics of the target of generalization.
2. *Ruling Out Irrelevancies*. Identifying those attributes of persons, settings, treatments, and outcome measures that are irrelevant because they do not change a generalization.
3. *Making Discriminations*. Identifying those attributes of persons, settings, treatments, or outcome measures that limit generalization.
4. *Interpolation and Extrapolation*. Generalizing by interpolating to unsampled values within a range of sampled persons, settings, treatments, and outcomes and by extrapolating beyond the sampled range.
5. *Causal Explanation*. Developing and testing explanatory theories about the target of generalization.

To demonstrate why Shadish's five principles are important, consider the following two sampling "realities" that researchers have to deal with. First, if questionnaires are sent to a random sample from a well-defined population, and the response rate is of the order of 30%, one simply cannot argue that the sample of respondents is random. To use "inferential statistics" or "statistical" generalization as Yin (2003) calls it, one is therefore forced to build a case for the representativeness of the sample, e.g. as discussed in Appendix 1. Second, researchers frequently seek to make general claims from case studies, experiments, or convenience samples (e.g. from surveys not based on probability sampling from a well-defined population). To use what Yin (2003) calls "analytic" generalization (discussed in Appendix 2), one must again build the case around claims of representativeness. In both situations Shadish et al.'s five principles are helpful for understanding what makes a sample (under which heading we include a single case study) representative of some broader class of things.

The intuition behind Shadish et al.'s principles is illustrated by the following example. Suppose that one has access to a convenience sample of data from 100 alumni from a business school at a Jesuit-backed¹¹ university in North America. Would knowledge claims based on analysis of that sample be representative of all business-school graduates in North America? The answer, we now show, depends on the proposition being generalized. If, say, the proposition concerns "behavioral intention to use a new technology", we argue that it is likely that the opinions, thinking, and patterns of behavior of these Jesuit-university-educated managers *would* be representative of North

¹⁰ Cronbach's 1982 book, *Designing Evaluations of Educational and Social Programs*, is an excellent source of well-informed views on the generalization of social-science research findings. For space reasons, a comprehensive summary of Cronbach's UTOS framework is not provided in this paper. However, use of this framework for discussing generalization is strongly recommended. Shadish et al. (2002) provide a useful summary of the framework.

¹¹ Christian Roman Catholic

American business school graduates, and probably of graduate managers, generally, in North America¹². However, if the proposition to be generalized relates to reasons for approval or otherwise of birth-control techniques, we argue that because of their Roman Catholic backgrounds, the opinions of these managers would not be representative of North American managers generally. In short, representativeness needs to be demonstrated *for the key factors related to the proposition being generalized*.

Shadish's five principles in Table 4 formalize the intuitive reasoning in the previous paragraph. Principle 1, *surface similarity*, is a simple initial test: Do things in the sample "look like" things in the population? In terms of the preceding alumni example—where the proposition of interest is intention to use a new technology—we judge that the managers in the convenience sample "look" more like graduates from other North American business schools than they "look" like, say, North American engineers or North American priests. However, this notion of "looking like" something is still too fuzzy. Which features should one look for and which features should one ignore? Principle 2, *ruling out irrelevancies*, tells us to look for differences in factors that matter *for the proposition being generalized*. In terms of the alumni example, religious beliefs are probably irrelevant for propositions about IT usage decisions. Principle 3, *making discriminations*, suggests that the researcher should identify boundary conditions beyond which generalizability claims do not apply. In the case of IT usage decisions, *interest* in technology is likely to be an important factor. Engineers, as a group, are probably more interested in IT than either MBA graduates or priests, and since intrinsic interest in the technology is likely to be an important factor determining intention to use, generalizability claims based on MBA graduates may not apply to engineers.

Principle 4, *interpolation and extrapolation*, involves using variance in the existing dataset to assess the sensitivity of research findings to variance in key concepts, and use this to predict variability in the instances not studied. For example, differences in responses between younger and older respondents *in the sample* may assist researchers in making judgments about applicability of their findings in various target populations. Finally, Principle 5, *causal explanation*, suggests that if cause and effect are believed to be understood, those cause-and-effect concepts are the ones to consider in assessing whether the sample is representative of the other settings of interest. In terms of our example, if the theory is that perceived usefulness and ease of use predict intention to use a new technology, then grounds for generalization are most soundly based on whether members of the sample and target population "think the same way" about usefulness and ease of use of IT. Graduates of other business schools in the US probably do. Engineers may. Priests in the US may not. Children in the US probably do not. And finally, Buddhist monks in Thailand almost certainly do not.

Future Discussions of Generalization in IS Research

Having (a) established that the state of generalizability practice in IS research is unsatisfactory, and (b) defined "other-settings generalization" and shown how representativeness may be demonstrated, in this third part of the paper we now ask: How should generalizability be discussed in future IS research reports? In other words, what should IS researchers say about the other-settings generalizability of their findings, and how much should they leave to their readers to make their decisions about the relevance of the findings to the reader's own situation (transferability)? In Table 5 we present three necessary elements for sound discussion of other-settings generalization in IS research papers.

Element 1 simply recognizes that since IS is an applied discipline, empirical IS research that does not discuss both generalizability and boundary conditions is logically incomplete. It also recognizes that discussing where claims might *not* apply may be very helpful to readers making transferability decisions.

Element 2 recognizes that generalization of knowledge claims from empirical studies relies solely on whether the subjects in the sample studied are representative of the target population. For sampling-logic research, e.g. surveys, discussion of other-settings generalizability would often involve simply trying to be more explicit about the sampling frame and population from which the sample was drawn, and revealing any concerns about non-response bias or choice of sampling frame causing lack of representativeness (King and He 2005). For small-sample studies, e.g. case studies or ethnographies, discussion of other-settings generalizability would often involve attempts to

¹² For this particular proposition, one way to convince readers that the sample is representative would be to compare proportions of males to females, proportions of managers of different ages, proportions in different discipline areas, and so on, with known proportions in the population of North American graduate managers.

define contexts where one would expect knowledge claims from the sample studied to also hold. IS researchers can operationalize Recommendation 2 by drawing on the vocabulary provided by Cronbach's (1982) UTOS or *UTOS framework and Shadish et al.'s (2002) five principles for generalizing in their discussions.

Element	Why this element matters
1: IS researchers should discuss explicitly the other-settings generalizability of their knowledge claims, including delineating clearly the boundaries beyond which their knowledge claims might not apply.	In our applied discipline, empirical research that does not discuss <i>both</i> generalizability and boundary conditions is logically incomplete. Table 2 shows that only 50% of the 70-odd papers where generalizability was relevant in ISR and MISQ in 2003 and 2004 discussed generalizability, and only 20% discussed boundary conditions. If these 70-odd papers are representative of the top IS literature—and logic says that they probably are—many IS papers are logically incomplete. The problem is that results based on an analysis of, say, IT usage in half a dozen US companies in 1999, do not necessarily apply in other contexts, e.g. in US homes, in US primary schools, in Sri Lankan businesses, or in Thai monasteries. Yet, without explicit statements to the contrary, readers may assume that research results do apply in settings where the authors would not expect their theories to be relevant.
2: Other-settings generalization claims should be based on the representativeness of the sample. This involves comparing key attributes of the sample and target populations, e.g. using Shadish et al.'s (2002) principles in Table 4.	Research that claims or implies generalizability to other settings, yet which does not establish sample representativeness <i>for the theory of interest</i> is logically unsound. For example, in survey research, if one's theory is about a relationship between x and y, one cannot establish sample representativeness by comparing attributes a and b for early and late respondents (unless a and b are logically related to x and y, respectively). Yet we can point to papers in our sample of 70 studies that compare early and late respondents on variables with no bearing on variables in the models tested. It may be harder to demonstrate representativeness for case-study research, but this does not mean it should not be attempted.
3: IS researchers should discuss known uncertainties or misgivings in all claims of other-settings generalization.	Given their own intimate knowledge of their data, research methods, etc., researchers <i>should</i> be in a better position than readers to understand the limitations and boundary conditions of their knowledge claims. For the same reasons that pharmaceutical companies include warnings about known adverse reactions to their drugs on their packages, so, good research reports should try to help readers—both practitioners and researchers—make sounder judgments about the transferability of research findings.

Element 3 follows from our scientific-realist truth definition which means that OSG claims are based on a reasonable expectation of wider applicability (which may subsequently be shown to be incorrect). It relies on the joint beliefs that (a) IS researchers have a *responsibility* to document known uncertainties and misgivings when making other-settings generalizations, and (b) some discussion of generalization is more useful than a blanket disclaimer of no generalizability, e.g. of case-based IS research.

In providing this advice, we are aware of the risk that authors and reviewers might honor the “letter of the law” in Table 5, and not the spirit of that advice. For instance, a researcher could simply claim that his or her results are not generalizable beyond the dataset studied. If so, the IS research community would be no worse off than at present. But our suspicion is that many researchers would welcome the opportunity to discuss openly their views about the likely generalizability of their findings to other settings. Perhaps some sort of “safe harbor” clause like that used when US firms present their financial statements (see the U.S. Private Securities Litigation Reform Act of 1995) would give researchers the freedom to say what they believe is likely to be true, rather than just that for which they have strong evidence. In the fullness of time, some claims of generalizability would, of course, prove to be wrong. This would provide valuable lessons for future researchers.

Applicability of the Arguments in this Paper to other Disciplines

Finally, in a paper advocating the inclusion of a discussion of other-settings generalizability at the end of all empirical IS studies, it would be remiss of us not to discuss the applicability of our arguments to other disciplines. Our position is that although our evidence of unsatisfactory generalization practice—from Lee and Baskerville (2003), King and He (2005), and our own analysis discussed in the first quarter of this paper—is limited to IS research, none of the key characteristics of other-settings generalization summarized in Table 1 is specific to IS research, so the recommendations in Table 1 *should* apply to any social-science research. (This is a Step-3-of-Figure 2 knowledge claim; it still needs to be tested.)

Conclusion

The contribution of this paper is the aggregation of ideas we have called “other-settings generalization” (OSG) and the recommendations that flow from this concept, presented in Table 5. OSG is arguably the most important subset of the range of concepts presented in Lee and Baskerville’s (2003) generalization framework. This set of ideas provides a clear, simple conceptualization of generalization that should be useful for future IS research, positivist or interpretive, quantitative or qualitative.

An other-settings generalization is the researcher’s act of arguing, based on the representativeness of the sample, that there is a reasonable expectation that a knowledge claim already believed to be true in one or more settings is also true in other clearly defined settings. As summarized in Table 1, OSG includes (a) recognition that all human knowledge is bounded, (b) recognition that all knowledge claims—including generalizations—are subject to revision, (c) an ontological assumption that objective reality exists, (d) a scientific-realist definition of truth, and (e) a distinction between generalization and transferability. Further, sound other-settings generalizations of research findings satisfy the following three requirements:

1. The researcher must clearly define the larger set of things to which the generalization applies;
2. The justification for making other-settings generalizations ultimately depends on the representativeness of the sample, not use of statistics;
3. Representativeness is judged by comparing key characteristics of the proposition being generalized in the sample and target population.

Having (a) shown that present practice is unsatisfactory, (b) presented our aggregation of ideas, we then (c) argued that practice would be much improved if IS researchers included in their empirically-oriented papers an explicit discussion of the other-settings generalizability of their findings, e.g. in the Discussion section of their papers. Such discussions, which could be framed in terms of Cronbach’s (1982) Units, Treatments, Observations and Settings (UTOS) framework and Shadish et al.’s (2002) five principles, would reveal assumptions and limitations of the study more effectively than is often the case at present. A particularly effective approach seems to be to focus on the causal drivers of whatever theory is in use in the paper (Principle 5), then to explore grounds for believing that those factors apply or do not apply in other settings. Although Shadish’s principles do not lead to black and white answers about other-settings generalizability, they are very helpful. Further, they clearly demonstrate that making sound other-settings generalization *always* involves human judgment.

Finally, we return again to the reasons why discussion of OSG is important in future IS research. First, IS is an applied discipline. It is therefore important that the implications for practitioners of IS research findings be discussed as clearly and openly as possible. Second, with increasing globalization of IS research, limits to generalizability of findings between different contexts need to be discussed more clearly than they have in the past. Third, since all knowledge claims are bounded, academic rigor demands that boundaries of applicability of any knowledge claims be carefully delineated. The three elements of a rigorous discussion of OSG are summarized in Table 5.

Appendix 1: Generalizing Using Inferential Statistics

Inferential statistics, or “statistical generalization” as Yin (2003) calls it, is concerned with drawing conclusions about a population based on observations of a probability sample selected (usually from a sampling frame) from that population. Any empirical study that reports a p-value (often indicated by notations such as “**”, $p < 0.05$) has generalized relying on this logic. However, not all empirical studies acknowledge that:

“The Achilles’ heel of statistical practice and inference is the selection of the sample. A probability model is *required* to draw inferences.” (van Belle, 2002, p.2) (emphasis added)

In practice, researchers rarely use random samples from well-defined populations (or sampling frames) because random samples are extremely difficult to collect. However, if the sample is not randomly selected, the probability theory that underpins inferential statistics is not directly useful for making generalizations. In the absence of random sampling, one has to resort to arguments about the *representativeness* of the sample to justify drawing conclusions about a population based on findings from a sample.

The inference logic when working with representative—as opposed to random—samples can be summed up as follows:

1. First one must establish that the sample is representative of some well-defined population. Such judgments are based on arguments similar to those from Shadish et al. (2002) presented in the body of the paper. They may well not involve statistics. For surveys, a common technique for attempting to establish representativeness is to compare early and late responses. What seems not to be well understood is that it is essential that variables *relevant to the study* are compared, not just any old factors that happen to be available.

2. Second, given a representative, but non-random, sample (e.g. responses from a typical survey where an author has presented evidence that the sample *is* representative of the target population), two arguments can be used for drawing conclusions about the population from which the sample was drawn.

Argument 1: Since the sample has been judged to be representative of the population, phenomena (e.g. means and path coefficients) observed in the sample correspond directly to phenomena in the population. In this case, just as when working with an entire population, no inferential statistical analysis is required.

However, since the estimate of population parameters is based on a sample, readers and reviewers will tend to ask questions about the confidence limits associated with such estimates. For example, Goodhue et al.’s (2006) Monte Carlo simulations show that sample-based estimates of path coefficients in their Table 1 do not always match those in the population. Therefore, most researchers pursue the second argument as well, or instead of, the first argument.

Argument 2: Since the sample has been judged to be representative of the population, it is equivalent to, or may be treated as, a random sample from the population¹³. One then applies the logic of inferential statistics to construct confidence intervals and estimate probabilities (p-values) that a null hypothesis is wrongly rejected. Although few studies make this clear, and few studies clearly define their assumed population, all published research studies that use non-random samples yet report p-values implicitly rely on this generalization logic. The p-values are indicators of the researchers’ confidence that null hypotheses can be rejected *in the population* from which the sample is assumed to have been drawn.

Conclusion on Representative Samples

This need to demonstrate representativeness of the sample (in terms of key criteria in the study) is rarely discussed in inferential-statistical analysis in published IS research. But for all except true random samples, it is a critical assumption in the logical chain that gives meaning to the computed p-values, t-statistics, etc. Even for true random samples, researchers would prefer representative random samples, because evidence of representativeness is grounds for believing that estimated population parameters are more likely to be within the desired (e.g. 95%) confidence limits. Thus, sample representativeness is a key requirement for generalization based on inferential statistics.

Appendix 2: Other-Settings Generalization and Yin’s Analytic Generalization

Yin (1989, 1994, 2003) is probably the most frequently cited authority on positivist case-study research. Through three editions, his book has proven very useful for countless case-study researchers. Yin argues strongly that

¹³ In fact, if it is truly representative on dimensions important to the study, distributional properties of the x’s and y’s and of the relationship between them will be more similar to those in the population than some random samples. This is because one deliberately sets out to exclude random samples (that might have been selected by chance) with distributional properties very different to the population.

generalization from case studies uses a *different* logic than that used in generalizing survey-based research. Specifically, he says that case study and experimental researchers “generalize to theory” using what he calls analytic generalization¹⁴, whereas survey researchers generalize from samples to populations using what he calls statistical generalization¹⁵. To help clarify the meaning of other-settings generalization, we now compare our OSG concept to Yin’s generalization concepts.

Our position is that we agree with Yin that theory generated from case studies may be validly generalized to other settings, and we agree that the logical pathways for analytic and statistical generalization are different, but we feel it is important to point out that both are other-settings generalization, and further, that the underlying logical basis for generalization in both logical pathways—the logic underlying the pathways—is the same in all other-settings generalization. In this Appendix, we demonstrate that

- (a) Both “analytic” and “statistical” generalization are simply examples of other-settings generalization,
- (b) Although the pathways or mechanisms for justifying analytic and statistical generalization are different, the underlying logical basis for all other-settings generalization relies on the *representativeness* of the sample.

Yin’s position on analytic and statistical generalization is stated most clearly in the following quotation:

“The external validity problem has been a major barrier in doing case studies. Critics typically state that single cases offer a poor basis for generalizing. However, such critics are implicitly contrasting the situation to survey research, in which a sample (if selected correctly) readily generalizes to a larger universe. *This analogy to samples and universes is incorrect when dealing with case studies.* Survey research relies on *statistical* generalization, whereas case studies (as with experiments) rely on *analytical* generalization. In analytical generalization, the investigator is striving to generalize a particular set of results to some broader theory” (Yin, 2003, p.37, emphasis in the original).

(a) “Analytic” and “statistical” generalization are both other-settings generalization

Although we agree with Yin that the logical pathways used in statistical inference and in generalizing from case studies to other settings are different, we argue that both are forms of other-settings generalization. With respect to samples and universes, the logic for statistical generalization is that, based on findings from a random or representative sample, assertions are made about relationships that probably exist in the population. This is other-settings generalization, namely, from the sample to the population. Turning now to case studies and experiments, the logical pathway for asserting broader applicability of the findings no longer uses statistics, but the result is the same. Based on arguments about the similarity of factors or subjects in the case or experiment and some other broader setting, the researcher argues that cause and effect relationships believed to exist in the case study or experiment will also apply in those other, broader settings.

For example, consider generalizing the findings from Kahneman and Tversky’s (1979) prospect-theory experiment. Kahneman and Tversky (1979) conducted their experiments with 72 subjects from Israel in the late 1970s, and discovered that their subjects’ attitudes toward risk aversion and risk seeking were different. In reporting their results, they seem to assume that these attitudes apply to all adult human beings, across all cultures, not just the subjects tested¹⁶. The only way to justify such a generalization is to argue that in making decisions such as those in the experiment, the sample of subjects is representative of all other people. The logical pathway for reaching this conclusion is not based on the probability calculations of statistical generalization (inferential statistics), but the consequence is the same. Based on findings from a sample of 72 subjects, conclusions are drawn or implied about a broader population of things, in this case “all people”. This, too, is other-settings generalization.

¹⁴ Level-2 inference in Yin 2003, Figure 2.2, p.32. This figure is reproduced in Lee and Baskerville (2003, p.222).

¹⁵ Level-1 inference in Yin 2003, Figure 2.2, p.32.

¹⁶ Kahneman and Tversky (1979) mention “essentially identical” results with “groups of students and faculty at the University of Stockholm and at the University of Michigan” (p.264-265). This helps give the reader more confidence in the generalizability of their results to other settings, but the limits of applicability are not discussed explicitly.

(b) The logic underlying other-settings generalization is representativeness of the sample

Although the mechanisms for justifying analytic and statistical generalization are different (quantitative studies may use inferential statistics, case studies do not), the logic underpinning both analytic and statistical generalization relies on the *representativeness* of the sample. In the case of analytic generalization, this is almost self-evident. For example, in Kahneman and Tversky's (1979) prospect-theory-experiments example above, the only grounds for generalizing from the finding with the 72 students is to assume that their attitudes toward risk aversion and risk seeking apply to all adult human beings, across all cultures. Likewise, if Robey and Sahey (1996) had wished to generalize their conclusion (based on two case studies of organizations using geographical information systems) that "information technology's consequences are socially constructed", the only logical grounds for so doing is to argue that the various social pressures that led to IT's consequences being socially constructed in their two case studies are also likely to exist in other settings, i.e., that relevant aspects of their case studies are representative of those other settings. In short, the grounds for analytic generalization clearly rest on the representativeness of the sample.

The grounds for believing that statistical generalization also relies on the *representativeness* of the sample have been presented in Appendix 1. In a nutshell, that argument is that while sample statistics such as regression coefficients and coefficients of determination (R^2) may be calculated from a sample without any consideration of the population from which the sample has been drawn, as soon as the researcher starts to consider tests of significance (e.g. t-tests, p-values) he or she has embarked into generalization territory. At this stage, the population must be clearly defined and the researcher must ask if the sample is (a) random, (b) representative, or (c) neither (e.g. a convenience sample). For other than random samples, representativeness must be demonstrated before one can argue that inferential statistics have meaning.

Summarizing, although the pathways or mechanisms for justifying analytic and statistical generalization are different, the underlying logical basis for both, and therefore for all other-settings generalization relies on the *representativeness* of the sample. Guidelines for deciding what factors to consider when assessing representativeness are presented in the discussion of Shadish et al.'s principles in the body of the paper.

References

- Baskerville, R. 1996. Deferring Generalizability: Four Classes of Generalization in Social Enquiry, *Scandinavian Journal of Information Systems* Vol.8
http://www.e-sjis.org/journal/volumes/volume08/articles/no2/01_baskerville_p5-28.pdf
- Bassellier, G., Benbasat, I., and Reich, B.H. 2003 The Influence of Business Managers' IT Competence on Championing IT, *Information Systems Research*, 14(4): 317–336
- Campbell, D. and Stanley, J. 1963. *Experimental and Quasi-Experimental Designs for Research*, Boston MA: Houghton Mifflin.
- Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design and analysis issues for field settings*. Chicago: Rand-McNally.
- Cronbach, L. J. (1982). *Designing evaluations of educational and social programs*. San Francisco: Jossey-Bass.
- Davis, F.D. 1989. Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology, *MIS Quarterly* 13(3): 319-339.
- Dillman, D.A. 2000. *Mail and Internet Surveys: The Tailored Design Method*. New York: Wiley
- Dubin, R. 1969. *Theory Building*. New York: The Free Press.
- Fan, M., Stallaert, J., and Whinston, A.B., 2003. Decentralized Mechanism Design for Supply Chain Organizations Using an Auction Market, *Information Systems Research*, 14(1): 1–22
- Fichman, R.G. 2004. Real Options and IT Platform Adoption: Implications for Theory and Practice, *Information Systems Research*, 15(2): 132-154.
- Goodhue, D., Lewis, W. and Thompson, R., 2006, PLS, Small Sample Size, and Statistical Power in MIS Research, *Proceedings of the 39th Hawaii International Conference on System Sciences*.
- Gregor, S., 2006, The Nature of Theory in Information Systems, forthcoming *MIS Quarterly*
- Hume, D., 1739. *A Treatise on Human Nature*, Reprinted in Hume, D., Norton, D.F., and Norton, M.J. (2000) *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*, Oxford University Press
- Hunt, S.D. 2002. *Foundations of Marketing Theory, Toward a General Theory of Marketing*, New York: M.E.Sharpe

- Hunt, S.D. 2003. Controversy in Marketing Theory, For Reason, Realism, Truth and Objectivity, New York: M.E.Sharpe
- Kahneman, D. and Tversky, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263-291.
- King, W.R. and He, J. 2005. External Validity in Survey Research, *Communications of the AIS*, Vol. 16: 880-894.
- Kirsch, L. 2004. Deploying Common Systems Globally: The Dynamics of Control, *Information Systems Research*, 15(4), pp. 374–395.
- Klein H.K., and Myers, M.D. 1999. A Set Of Principles For Conducting and Evaluating Interpretive Field Studies In Information Systems *MIS Quarterly* Vol. 23 No. 1: 67-94.
- Kohli, R. and Kettinger, W.J. 2004. Informing the Clan: Controlling Physicians’ Costs and Outcomes, *MIS Quarterly*, 28(3): 363-394
- Lee, A.S. and Baskerville, R.L. 2003. Generalizing generalizability in Information Systems Research. *Information Systems Research* 14(3): 221-243.
- Leidner, D.E. and Kayworth, T. 2006, A Review of Culture in Information Systems Research: Towards a Theory of IT-Culture Conflict, forthcoming (June), *MIS Quarterly*.
- McGuire, W.J. 1983. A contextualist theory of knowledge: Its implications for innovation and reform in psychological research. In L. Berkowitz (Ed.) *Advances in experimental social psychology*, (Volume 16, pp. 1-47). New York: Academic Press.
- Melville, N., Kraemer, K., and Gurbaxani, V. 2004. Review: Information Technology and Organizational Performance: An Integrative Model of IT Business Value, *MIS Quarterly*, 28(2): 283-322.
- Monteiro, E., and Hanseth, O. 1996. Social Shaping of Information Infrastructure: On Being Specific about the Technology, in *Information Technology and Changes in Organizational Work*, W. J. Orlikowski, G. Walsham, M. R. Jones and J. I. DeGross (eds.), Chapman and Hall, London: 325-343
- Neuman, W.L., 2003. *Social Research Methods: Qualitative and Quantitative Approaches*, 5th ed. Boston: Allyn and Bacon.
- Palmquist, M. et al. 2004. “Introduction to Generalizability and Transferability”. Website: <http://writing.colostate.edu/guides/research/gentrans/pop2a.cfm> (viewed May 2006)
- Popper, K.R., 1992. *The Logic of Scientific Discovery*, 5th ed, London: Routledge
- Raghu, T.S., Jayaraman, B. and Rao, H.R., 2004. Toward an Integration of Agent- and Activity-Centric Approaches in Organizational Process Modeling: Incorporating Incentive Mechanisms, *Information Systems Research* 15(4): 316–335.
- 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.
- Shadish, W.R. 1995. The logic of generalization: five principles common to experiments and ethnographies, *American Journal of Community Psychology*, Vol 23 #3: 419-428.
- Shadish, W.R., Cook, T.D., and Campbell, D.T. 2002. *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Boston MA: Houghton Mifflin Company
- van Belle, G. 2002. *Statistical Rules of Thumb*, Wiley
- van der Aalst, W. M. P. and Kumar, A. 2003. XML-Based Schema Definition for Support of Interorganizational Workflow, *Information Systems Research* 14(1): 23–46.
- Venkatesh, V., M.G. Morris, F.D. Davis, and G.B. Davis, 2003. User Acceptance of Information Technology: Toward a Unified View, *MIS Quarterly*, 27 (3): 425-478.
- Walsham, G. 1995. Interpretive case studies in IS research: nature and method, *European Journal of Information Systems*, 4(2): 74-81.
- Walsham, G. and Waema, T. 1994. Information Systems Strategy and Implementation: A Case Study of a Building Society, *ACM Transactions on Information Systems* (12:2): 150-173.
- Weber, R. 2004. Editor’s Comments—The Rhetoric of Positivism Versus Interpretivism: A Personal View, *MIS Quarterly* (28,1) March 2004: iii-xii.
- Weick, K. E. 1985. Theoretical assumptions and research methodology selection. In McFarlan, F. W. (Ed.), *The information systems research challenge* (pp. 111-132). Boston, Massachusetts: Harvard Business School Press.
- Whetten, D.A. 1989. What constitutes a theoretical contribution? *Academy of Management Review* (14,4): 490-495.
- Wood, A. 2000. *Hume: The Problem of Induction*. Stanford University. <http://www.stanford.edu/~allenw/Phil102/Hume%20-%20Induction.doc>. (viewed August 1, 2006)
- Yin, R.K. 1989, 1994, 2003. *Case Study Research, Design and Methods*, 1st, 2nd, and 3rd editions, California: Sage Publications.