7-29-2001

Conducting Experimental Research in Information Systems

Alan R. Dennis
Indiana University, ardennis@indiana.edu

Joseph S. Valacich
Washington State University, jsv@wsu.edu

CONDUCTING RESEARCH IN INFORMATION SYSTEMS

Alan R. Dennis
Kelley School of Business
Indiana University

ardennis@indiana.edu

Joseph S. Valacich
College of Business and Economics
Washington State University

TUTORIAL

CONDUCTING RESEARCH IN INFORMATION SYSTEMS

Alan R. Dennis
Kelley School of Business
Indiana University

ardennis@indiana.edu

Joseph S. Valacich
College of Business and Economics
Washington State University

ABSTRACT

This article presents a summary of key success factors for publishing research in top-tier IS journals; it is not intended to be an introduction to research, but to go beyond the "rational model" presented by most introductory works. The paper begins by discussing the processes by which research projects are identified and developed, specifically focusing on where project ideas are found and how projects are selected and refined. Next, we discuss the fundamental role that theory development, testing and refinement plays in research. This discussion is followed by an examination of several interrelated research design issues, including maximizing publication potential, and executing the study's activities. Next, the importance of writing quality as well as the cultivation and refinement of a project's message is discussed. Finally, a checklist is provided on "how to be rejected" which summarizes the central themes of this article.

Keywords: research methods, research design

I. INTRODUCTION

Many good books provide useful introductions to the design and execution of research, particularly laboratory experiments (e.g., Kerlinger, 1986; Babbie, 1995). These books are important tools for beginning scholars because they provide a solid, logical presentation of the key issues in research design. However, after publishing several dozen research papers, we believe that while having an understanding of the fundamental issues presented in these basic texts is necessary, there are other requirements necessary for the publication of research in top-tier information systems (IS) journals.

Most of these methodology books were written by authors outside IS and outside Schools of Business. While most principles of research design transcend disciplines, we found from personal experience in publishing in IS, management, speech communication, and psychology, that different disciplines develop different norms for the design and presentation of research. Based on our experience as authors, reviewers, and associate editors, we can honestly say that the leading IS journals today are at least as rigorous as the best journals in other disciplines. Most of these basic texts on research design present what we would call the "rational model" of science. Many of what we believe are the key success factors in publishing excellent research are not found in the "rational model" that often dominates what we teach our beginning scholars.

The goal of this article is to present a summary of what we believe are the key success factors in publishing research in top tier IS journals. This paper is not a basic introduction to research, because, as noted above, many good introductions to research are far more ambitious in scope than this article. We assume that the reader has a solid understanding of the basics of research design. Our focus is on the next step: how to move beyond the basics to become an expert practitioner; that is, the "street smarts" that come from practice (cf. Jarvenpaa, Dickson and DeSanctis, 1985).

In this article we focus on four aspects of research design, using one study as an example of the principles we develop.

- First, we briefly present our view on the role, and the inherent tradeoffs, of choosing a research methodology (Section II).
- Next, we discuss the development of research projects: the selection of a research team and the identification of important research questions (Section III).
- Third, we examine the role of theory in research design (Section IV).
- Fourth, we provide practical guidelines for research design aimed at improving the probability of producing a top tier article (Section V).

We then discuss key elements in writing up the study for publication (Section VI). Finally, we conclude with our list of the top ways to have your paper rejected. Throughout this article we will use a series of examples to demonstrate and elaborate on the key points we raise.

II. THE ROLE OF RESEARCHMETHODOLOGY

Many of our examples and much of our discussion in this article focuses on the design of quantitative research for the simple reason that most of our work has been quantitative. Quantitative research methods are usually the tools of researchers who examine phenomena from a positivist perspective. Researchers who adopt a positivist perspective assume "the existence of a priori fixed relationships within phenomena (p. 5) … whose nature can be relatively unproblematically apprehended, characterized, and measured" (p. 9) (Orlikowski & Baroudi, 1991). Our goal is not to debate the merits of positivism, or to advocate any epistemological perspective (see, however, Orlikowski & Baroudi, 1991; Lee, 1991). We firmly believe that both quantitative and qualitative research are essential elements in IS research, but since we lack rich experience with qualitative research this article will focus primarily on quantitative research.

We believe that the purpose of research is to advance knowledge and the scientific process. Each scientific method has its strengths, and unfortunately, all methods of science are flawed. For example, one of the most common criticisms leveled at experimental research is that it is artificial. How can studying undergraduate students working on pretend tasks for which they have little

interest or experience possibly be "real"? Isn't laboratory research seriously flawed? Well, yes. The critics are right: all laboratory experiments are seriously flawed. However, *all* research methods are seriously flawed. One of the best discussions of the limitations of experimental research -- and survey research and field research -- is that by McGrath (1982). McGrath (1982) argues that research methods can be evaluated on three dimensions:

- *Generalizability* with respect to populations
- *Realism* for the participants
- *Precision* in the control and measurement of variables.

It is literally impossible to design a research study that satisfies all three dimensions, although sometimes it is possible to strike an uneasy balance among two of the three (and fail miserably on the third) (McGrath, 1982). Laboratory experiments, for example, maximize precision, but usually fail to satisfy generalizability or realism. Field studies maximize realism, but fail to satisfy generalizability (because they study a small number of non-randomly selected situations) or precision (because there are a host of uncontrolled factors). Surveys maximize generalizability, but fail to satisfy realism (because they do not study actual behavior but instead ask participants to recall perceptions) or precision (because there are a host of uncontrolled factors).

Because all research methods are imperfect, anyone claiming that experimental research is too seriously flawed to be used or that surveys or field studies are better is simply ignorant. No one method is better or worse than any other; they are simply better at some aspects and worse at others. Therefore, to truly understand a given phenomena, we believe it is important to study it using different methods across a series of different studies. Under ideal circumstances, it may even be possible to combine experimental research with qualitative analyses drawn from field research to understand the issues better (e.g., Trauth and Jessup, 2000; Lee, 1991).

In summary, we do not propose that quantitative methods are the "only" or the "best" research method, but we do propose that they are useful and necessary in a community of scientists that share the goal of gaining a deeper

and more comprehensive understanding of IS phenomena. We agree with Kurt Lewin that "There is nothing as useful as a good theory."

III. DEVELOPING RESEARCH PROJECTS

IDENTIFYING PROJECT IDEAS

The first and most important aspect of any research project is to develop the research team and the key question(s) the project will address. From a "rational model" standpoint, there is a logical progression. The project starts with previous research and theory, which leads to research question(s), which in turn drives the choice of research methods and ultimately the design of the research study (Martin, 1982).

In our experience, research is never that orderly. We believe the "garbage can model of research" (Martin, 1982) is a more useful model of how research projects are typically developed. The garbage can model argues that there is no orderly progression; instead the key elements of the project are thrown together into a garbage can, mixed together, and out comes the project.

In some cases, project ideas do come from prior research and theory, as the rational model argues, but, as shown in Figure 1, in most cases ideas come from a much wider and richer set of stimuli than the rational model would have one believe (Martin, 1982). Some of the most powerful projects come from personal experiences, particularly when those experiences clash with the dominant wisdom from prior research and theory. Our earliest research projects on electronic brainstorming (Nunamaker, et al. 1991) were prompted by a mismatch between published research and theory, and what we were observing in the field. Our qualitative field research suggested that Group Support Systems (GSS) use could improve performance, but much published research at the time (all of which was quantitative) suggested that GSS use did not improve performance.
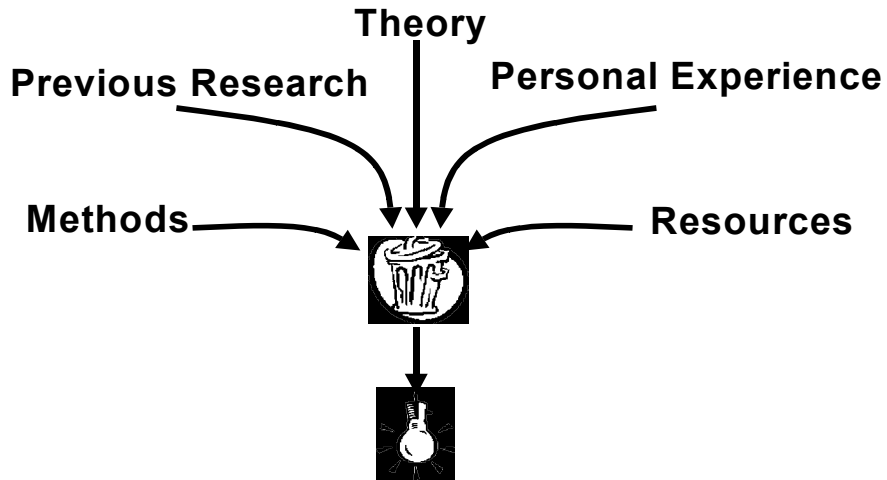
Figure 1. Sources of Research Ideas

Methods and available resources also play a key role. Our first decision on the very first research project we did together was that we wanted to do a laboratory experiment on GSS performance; the choice of method came before the choice of the research question. In our case, we had a wonderful facility that enabled us to do laboratory research and a large subject pool to draw upon. Without the subjects and the facility (the resources), the project would not have been possible. Thus the idea for the project came from a mix in the garbage can, not from a well structured "rational" progression.

SELECTING A PROJECT

With so many sources of ideas mixing together, the problem is usually not trying to generate an idea for a project, but rather choosing among ideas and refining them into a workable project. So how do you identify and select a good research project? Obviously, the project should fit the interests of the individuals involved in the project. It should be fun and interesting, or the lack of motivation will ultimately doom the project. Assuming that the projects are interesting to those poised to undertake them, the next three subsections convey important considerations in selecting and refining research projects.

Publication Potential

The first consideration is the potential for publication. In our opinion, a good research project is one that has the potential to be published in a top-tier journal. The goal of all research is to publish a new contribution to knowledge -- something that extends to the cumulative body of knowledge so that we as a discipline now know something new. However, the goal of those striving to publish in top tier journals is not to be published, but rather to be read, to be cited, and to change practice.

This goal is what differentiates a study published in an top tier journal from that published a lower tier journal. Top tier journals strive to publish research that will be read, cited, and used, both by other researchers in their research, and eventually by practitioners. Very early in our careers, a senior editor of a top journal explained this concept to us by saying "your goal is to produce timeless classics." While not every research project will produce a timeless classic, we believe there are four questions that can help researchers understand the potential for producing timeless classics.

1. *Fundamental Issue.* Does the project address a fundamental issue, an issue that is likely to still be of concern in five years time? Top tier journals strive to publish archival research that continues to provide value long after the article is published. This is one reason, for example, that so little academic research was devoted to the Y2K problem.

2. *News Value.* Does the project have news value? That is, is there something clearly "new" about the project, in that it asks a totally new research question, or asks an old question in a new situation or in a new way that is likely to produce a different answer?

3. *Interesting story.* Does the project have an interesting *story* that is likely to appeal to reviewers and editors?

4. *Outcome independence.* Can the project make a contribution regardless of the outcome? For example, a recent experiment by Stephen Hayne and colleagues (Hayne and Rice, 1997) investigates whether participants in a groupware meeting can accurately identify the authors of anonymous

comments. In this case, the experiment has the potential to make a contribution whether it finds that participants can or cannot make accurate attributions of authorship; that is, regardless of whether there is a statistically significant effect.

If you identify a project where the answers to these questions is "yes", we believe the project has the potential to become a timeless classic.

Fit with Current and Future Research

The project should either be a first step or a next step in a stream of research. Research is simpler when individual research projects are related. With related studies, the core theory and body of previous research changes only slightly from study to study, and much of the prior work can be reused in developing the theoretical arguments and experimental materials for the next study. The goal is to continue to work a line of research until it is played out (Watson, Satzinger, and Singh, 1994).

When we do research, we often plan two or more separate studies at a time so that we can develop theory, measures, and materials for the two studies at the same time. Often the data collection effort is done at the same time as well. By combining several studies into one plan, one can become more efficient at doing research.

Project Risk and Return

Finally, the project should fit the risk/return portfolio of the individuals involved. The most precious resource available to any faculty member is research time. As with financial investing, investing time in research projects should be done with an eye to both the possible return from a successful project as well as the risks of project failure. Publications in top tier journals usually bring a greater return than publications in lower tier journals. However, they often are higher risk projects, because top tier journals require a more substantial contribution before publication. Gordon Davis (1992) provides an excellent methodology for assessing project risks that includes the assessment of completion risk (how likely you are to complete the project), opportunity cost risk (how much time the project will consume and limit your ability to undertake other

projects), publication risk (how likely the project is to produce a publication) and competition risk (how likely it is that someone else is working on a similar project).

In our opinion, the best publication strategy for junior and mid-career scholars is one that contains a balanced portfolio of high risk/high return projects and low risk/low return projects. In practical terms, this means that some projects should be aimed at top tier journals that have a higher risk of rejection but a much greater payoff once they are published, while other projects should be aimed at lower tier journals with a lower risk of rejection and a lower payoff upon publication. Promotion to full professor should bring with it a shift in focus. While full professors will need to continue to mentor junior scholars across a range of projects, we believe that senior scholars, freed from time pressures associated with achieving tenure and promotion, should strive to invest their time to a greater extent in higher risk/higher return projects.

RECOGNIZING POTENTIAL

Recognizing the potential of different projects can be challenging, especially for junior scholars. One of the most important and simple steps is to solicit the opinions of senior scholars about the potential of research projects before they are undertaken. As with software development projects, the simplest time to make major changes is at the concept stage.

Many famous scientists have discussed how they began projects that ultimately led to crucial scientific discoveries. In many cases, the individuals began projects because of hard to explain "gut" instincts. However, two dictums from the natural sciences seem particularly appropriate for guiding IS research (we wish we could provide accurate citations of the original sources, but these owe as much to folklore as a direct quote):

- The definition of a Nobel Prize is physics is "Oh #&@%, why didn't I think of that?" The key message is to seek simple solutions to complex problems.
- Scientific discovery does not begin with the word "Eureka;" it begins with the words "That's funny." The key message is to seek solutions to anomalies and paradoxes.

BUILDING A RESEARCH TEAM

Many years ago, much research was done by individual researchers working by themselves. Today, most research is done by teams of researchers. While individuals are still capable of doing research, most researchers recognize the value of working in teams. Most research requires three distinct skills.

1. The researcher(s) must be able to conceptualize and develop theory and hypotheses for the phenomenon under study. It may be better, for example, to build a cross-discipline research team to develop a more integrative theory.

2. The researcher(s) must be able to design and execute a robust research study to test the hypotheses.

3. the researcher(s) must be able to write a convincing paper that reports the theory and results.

Some of our colleagues would argue that a fourth element is also needed: the time to collect and analyze the data.

While some researchers are experts in all three areas, most of us are not so fortunate. The purpose for building a research team is to draw together individuals who provide complementary skills (or in the case of novice scholars such as junior Ph.D. students, time to collect and analyze the data).

AN EXAMPLE

In this article , we will use the experiment by Dennis and Kinney (1998) as an example to illustrate how to implement the principles we develop. We should be careful to point out that we are not claiming that this article is a perfect exemplar of ideal IS research; all studies are flawed and this one is no exception.

This experiment tested Media Richness Theory (MRT) (Daft and Lengel, 1986). MRT argues that media differ in "richness," defined as "the ability of information to change understanding within a time period" (Daft and Lengel, 1986, p. 560), depending upon the multiplicity of cues they can transmit (e.g., vocal inflection, gestures) and the immediacy of feedback they provide (i.e., how

quickly can a receiver respond to the sender)[1]. MRT further argued that performance would be improved when richer media were used for equivocal tasks (where there are multiple and possibly conflicting interpretations for information) and when leaner media were used for non-equivocal tasks.

The experiment used a 2 x 2 x 2 repeated measures design in which 2-member groups performed both an equivocal task (a university admissions task) and a less equivocal task (a standardized multiple choice college entrance exam task) using a medium that either provided high multiplicity of cues (video-conference) or low multiplicity of cues (text-based messaging) and either immediate feedback (subjects could send and receive at the same time) or non-immediate feedback (only one subject could send at a time and the other could not interrupt until the sender relinquished control).

Table 1 uses the Dennis and Kinney (1998) experiment to illustrate the key issues in developing the research project. The project began in January of 1992 when Susan Kinney, then a Ph.D. student, approached Alan Dennis, then a new assistant professor, with the idea of doing a joint project. Kinney was in the final stages of her dissertation and wanted to do another lab experiment using the subjects in the introductory MIS course. She had just finished analyzing her dissertation data that used MRT (Kinney and Watson, 1992) and was surprised that MRT "didn't seem to work." Dennis, having just completed his Ph.D. at the University of Arizona with its focus on electronic communication and groupware, believed that face-to-face interaction might be overvalued for work performance. Reflecting on her own experiences, Kinney realized that e-mail was a useful tool, but valued the personal contact of face-to-face interaction.

At this point, both agreed that a follow-up experiment to Kinney's dissertation was called for. But what exactly should it be? In reading the prior MRT research, both realized that empirical research had deviated from the theory's original intent as a theory of media use (i.e., predicting performance from

---

[1] Daft and Lengel also argued that richness depended upon the ability to personalize the message (e.g., a memo written to one person can be tailored to a greater extent than an flyer sent to hundreds) and the language variety the medium could support. These constructs were not included in the Dennis and Kinney experiment, so we omit them from our discussion.

Table 1. Developing the Research Project in Dennis and Kinney (1998)

Sources of Project Idea	• **Method:** Desire to do an experiment. • **Resources:** 300 student subjects per semester, computer lab available for use. • **Personal Experiences:** Prior experience with MRT and media that did not follow the conventional wisdom.
Selecting a Project	**Publication Potential** • MRT was a *fundamental issue* (media richness is still an issue today). • The study was *newsworthy* because no experiment had tested performance by varying media characteristics and task equivocality. • The study was likely to have an *interesting story* by showing MRT "worked" as a theory of performance while other research showed it didn't "work" as a theory of media choice. • Whether MRT was or was not supported, the study would *still have value*. **Research Fit** • Kinney had just completed an experiment using MRT the literature was fresh in her mind and many of the experimental materials could be reused. • Dennis was working in groupware and MRT was one of the basic theories in the area. **Risk and Return** • *Completion risk* was seen as low because both researchers had successfully completed prior related experiments. • *Opportunity cost risk* was seen as low because experimental materials existed and the research literature was known. • *Publication risk* was seen as low because of the publication potential above. • *Competition risk* was erroneously seen as low as no other researchers had to then used experimental research to investigate MRT.
Recognizing Potential	• Anomalies in published research (a theory of media use being tested as a theory of media choice) and y between the theory and the researchers own use of electronic media
Building a Research Team	• Dennis had conducted and published several lab experiments. Kinney was familiar with MRT, had extra experimental materials, and time to collect the data.

use) to a theory of media choice (i.e., why managers choose media). All research testing MRT had been field or survey, so precision was a major issue over the body of research. Most research had concluded that MRT was not ideal in predicting choice, but virtually all research to that time had simply accepted that the performance predictions of MRT were true; the theory itself had not been tested because no study had examined performance under different media and task combinations. Thus, as shown in Table 1, there was significant publication potential in testing the theory itself.

The project fit with the then-current research streams of both Kinney and Dennis. Few risks were foreseen, but unbeknownst to Kinney and Dennis, Joe

Valacich and a team of his doctoral students were independently undertaking a nearly identical study (Valacich, et al., 1994; Mennecke, Valacich and Wheeler, forthcoming). Ironically, this was to be the first in a series of three occasions that Dennis and Valacich independently and unknowingly conducted almost similar experiments. It just shows how the principles for selecting projects can lead to similar results when the contents of the garbage can generating ideas are similar at different places.

IV. DEVELOPING THEORY

A crisp, clear, concise theory is essential for all quantitative research -- and we would argue also for qualitative research, although the theory in qualitative research often comes from the data, whereas the theory comes before the data in quantitative research. It is often said that quantitative research is a theory in search of data, while qualitative research is data in search of a theory. The sole reason for publishing an article in a top tier journal is for the theory it develops; empirical data are useful only for supporting the theory, or rejecting it and opening the door to new theoretical advances. Many papers discuss how to develop and write good theory (e.g., see Bailey, 1987; Bacharach, 1989; Sutton and Staw, 1995; Van de Ven, 1989; Weick, 1995; Whetten, 1989). Rather than repeat these good works, we will focus on those aspects we see as most problematic in developing theory.

Theory is the *why* of the phenomenon, not the *what*. Theory explains the key actors in the phenomenon under study (the independent and dependent variables), how they interact (the plot), and why they interact as they do (their motivation). In the same way that a book or movie would be uninteresting if we did not understand the characters' motivation, so too is a research study that lacks theory. Theory is what enables us to generalize the results of a study beyond its boundaries; when we generalize, we generalize from the theory not from the data.

Theory is any set of logical arguments that explain a relationship among a set of constructs. In most cases, the theory for a particular study will build on and

instantiate a prior theory, modified as needed, or integrate a set of theories within the bounds of the study. In other cases, the study will build new theory based on prior empirical research and logical argument.

There are two key points here. First, theory is not a summary of prior research. Theory can summarize prior research, but must go beyond the empirical data to explain why the data are the way they are. "Data describe which empirical patterns were observed and theory explains why empirical patterns were observed or are expected to be observed." (Sutton and Staw, 1995, p. 374, original emphasis). Second, theory does not need to be a "Big T" theory. A "Big T" theory is an overarching theory that is widely recognized and used, such as Media Richness Theory (Daft and Lengel, 1986) or Adaptive Structuration Theory (DeSanctis and Poole, 1994). "Big T" theories often have formal names (written or abbreviated in capital letters, of course).

In contrast, "little t" theory is more narrowly focused within a more constrained space, such as theory on electronic brainstorming (Valacich and Dennis,1994). "Little t" theories are often more useful within their smaller domain because they are more focused. In our experience, many junior scholars often feel a burden to develop a "Big T" theory -- and often shy away from the prospect -- when a "little t" theory is simpler and more appropriate. At a minimum, "little t" theory is simply a set of boxes (constructs) and lines (relationships) with clear explanations of the nature of the relationships, why they exist, and the boundary conditions.

In the rational model of science, writing theory usually moves in a linear fashion from constructs, to relationships, to hypotheses. In our experience, writing theory is like developing an information system. Parts of the theory are built using the "top-down" rational science model and the "middle out" approach of starting with the relationships and seeing what constructs lie at the ends. In any event, much iteration and prototyping is required.

In the case of Dennis and Kinney (1998), the project focused on a "Big T" theory, so much of the discussion above does not apply. However, the

theoretical argument used to support the theory evolved as the paper moved through the review process, as will be discussed later.

Another example where "Big T" theory was the theoretical foundation for a study is reported in Wheeler and Valacich (1996). In this study, a specific instantiation of Adaptive Structuration Theory (AST) is used to investigate three distinct appropriation mediators -- facilitation, GSS configuration, or user training -- on how groups adopted technology, and how this use influenced decision performance. Since AST does not make predictions related to decision performance, "small t" theory was blended with the non-deterministic aspects AST to make specific predictions on decision quality when various mediators were introduced. This blend of "Big T" and "little t" theory was referred to as PRAST (Process Restricted AST). Thus, it may be appropriate to blend both "Big T" and "little t" theory within a single study.

Finally, in other cases, we have simply developed "small t" theory focused on the interplay of interesting independent variables. For example, Connolly, Jessup and Valacich (1990) examined the interplay of anonymity and the evaluative tone of a confederate group member who would inject supportive or critical comments into a group brainstorming session. To make predictions related to idea quantity and quality, member satisfaction and so on, several "small t" theories were developed (referred to as the "balance of forces model"), with no overarching "Big T" theory motivating the work.

In summary, strong theorizing is a fundamental part of the research process. In some cases, adopting a "Big T" theory might be appropriate, while in others "small t" or some combination of both "Big T" and "small t" might be appropriate. There is no one or correct approach other than the guideline that the only reason for doing research is to develop, test, and extend theory. Put simply, without theory, there is no science.

V. DESIGNING QUANTITATIVE RESEARCH

Many elements of a research study can be improved after the study is conducted. The theory can be revised, the statistical analyses redone, and the

paper rewritten. The one element that cannot be reworked is the research design itself. Once you collect the data, you are committed. Therefore, it is essential to ensure the basic research design is sound.

UNDERSTANDING THE STRENGTHS AND LIMITATIONS OF RESEARCH METHODS

As we argued above, all research is flawed. The best research designs, regardless of method, are those that openly accept their flaws and aggressively play to their strengths. *The primary strength of experimental research is precision and control; its primary purpose is to test and extend theory*. Experimentation is not intended to produce generalizable results -- although generalization can be achieved after an accumulation of studies that vary aspects of the research design. The important implication is that in designing an experiment, one should consider realism and generalizability, but focus on precision above everything else. If there are any compromises that must be made, it is important to sacrifice realism and generalizability for precision, because precision is the raison d'être of experimental research. The challenge, of course, lies in doing precise experimental research that is also interesting and relevant to the academic community at large.

Likewise, the primary strength of survey research is generalizability. Surveys are not intended to be realistic or precise -- although these too can be achieved after an accumulation of studies. Identifying the appropriate population that is representative of that under study is critical because that is raison d'être of survey research. The primary strength of field research is realism, so concerns about precision and generalizability are secondary to obtaining realism.

MAXIMIZING PUBLICATION POTENTIAL

Maintaining Construct Integrity

An important element of the research design of all studies is to maintain the integrity of the constructs, both the independent constructs (treatments, in the case of experimental research) and the dependent constructs (measured variables). Measurement validity and reliability is discussed at length in most

introductory research design texts (e.g., Babbie, 1995; Cook and Campbell, 1979; Kerlinger, 1986). However, in our experience validity and reliability are followed less often than most scholars would admit.

In our experience, validity is in the eye of the beholder, so provided the measures are reasonable in the eyes of the reviewers and the editor, validity is less of an issue than reliability. One simple approach is to reuse measures that were used and found valid in prior research. To test reliability, all dependent variables must use several items to measure them. Measures gathered via questionnaires should have at least three (ideally five) questions whose reliability is assessed via Cronbach alpha. Measures gathered via coding (e.g., counts of the number of ideas, assessment of solution quality) must have at least two separate individuals code them, whose inter-rater reliability is assessed via a simple percentage of agreement or via Cohen's kappa. To save coding time, it is acceptable to examine inter-rater reliability only over some reasonable subset of the entire data set (e.g., one third).

Testing the Theoretical Linkages

Since the primary reason for publishing experimental research is for the contribution to theory, the more evidence that you can provide to support or refute the theory, the stronger the study. One of the key elements in any theory is the explanation of why a relationship exists. When designing a laboratory experiment, it is useful to test the factors that underlie relationships. For example, suppose that we theorize that the use of electronic brainstorming should increase the number of ideas produced compared to verbal brainstorming for a number of theoretical reasons, including the fact that we believe that the anonymity in the electronic system will reduce the apprehensiveness about suggesting a "silly" idea. We obviously need to test the end conditions (i.e., number of ideas). However, we also should test the intervening relationships; that is, did the participants in the electronic treatment feel they were more anonymous than those in the verbal treatment and did they feel less apprehensive about contributing comments.

If the end-to-end conditions support the theory (i.e., the electronic groups did generate more ideas) then a test of the intervening variables that also supports the theory provides a more solid case, and helps us to understand the theory better. It also may help us understand which underlying factors are more important because if one theoretical linkage is not supported but another is, we know far more than if we had not tested the underlying linkages (e.g., if participants did not feel less apprehensive, then we might conclude that anonymity was less important than other factors in influencing the production of more ideas).

The tests of the underlying linkages become much more important when an end-to-end test of the theory fails. That is, the results do not support the theory. In this case examining the underlying linkages enables one to see exactly where the theory fails and helps us understand how we can repair the theory to make it useful (e.g., Dennis, Hilmer, and Taylor, 1998).

Designing for Statistical Significance

The goal of research is to support or refute a theory. For quantitative research, the ultimate test of the theory lies in the statistical analyses of the empirical data that are collected, so it is crucial that the research design be developed to test the hypotheses efficiently with the most statistical power that can be brought to bear. Therefore, quantitative research should be designed with statistical power in mind at all times. Every research design decision should be examined to see the implications it has for statistical power (Baroudi and Orlikowski, 1989).

The statistical heart of most quantitative research designs is the t-test, or one of its cousins such as the F-test, ANOVA, regression, and Lisrel. Figure 2 presents the simple equation for the t-test that should be second nature for most researchers. To reject the null hypothesis (the null hypothesis assumes that the proposed theory is incorrect), the t-statistic must exceed some level that is beyond a reasonable doubt, which is usually set at an alpha level of .05. The goal of experimental design is to adjust the elements within this equation so that there is the greatest a priori chance of rejecting the null hypothesis if it is false;

that is, to increase the value of t within the rules of experimental design (Kerlinger, 1986). If the null hypothesis is true (and the underlying theory is false), then there is nothing that you as a researcher can do to change the outcome -- barring, of course, unethical behavior.

The equation for the t-test (Figure 2) includes three parts that can be systematically adjusted through the research design (Kerlinger, 1986). The first is numerator, which is the difference between the treatment means ($X_1 - X_2$). To increase the value of t, the goal is to increase the potential differences between the treatment means in an experiment, or the independent variables in a survey. For experiments, the treatments or should be designed to differ as much as possible on the manipulated variable; subtlety is not desirable. Furthermore, there should be as few treatments as possible; if there are several possible levels for a manipulated construct there should be as few levels as is meaningful and they should be as noticeably different as possible. For surveys, select a population in which the independent variables are likely to present a wide variation in the sample population. Take care, however, not to make the experimental manipulations so large or the survey population so disparate that the results become obvious and uninteresting. Your goal is to walk a fine line: design to maximize the difference between the independent variables, but at the same time making sure that the differences aren't so great that they become uninteresting.

The second element of the t-test formula in Figure 2 is the standard deviation(s). To increase the value of t, the goal is to decrease the standard deviation relative to the sizes of the means, which can be addressed in three ways.

1. Attempt to use a population whose members are as homogenous as possible on the dependent variable. (Homogeneity is the reason than many medical trials are only run on men between the ages of 50 and 55).

Conducting Research in Information Systems by A.R. Dennis and J.S. Valacich

$$t = \frac{\boxed{X_1 - X_2}}{\boxed{s} / \sqrt{\boxed{n}}}$$

Maximize the difference of the treatment means

Increase the sample size
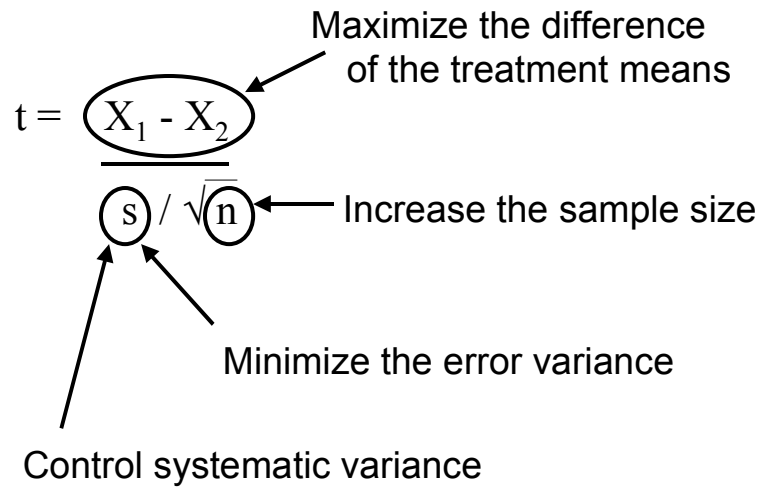
Minimize the error variance

Control systematic variance

Figure 2. Designing for Statistical Significance

2. Attempt to control any systematic variance in the population though the use of covariates. For example, if one is studying the effects of different idea generation techniques on the number of ideas produced, then it might make sense to measure each subject's innate creativity as a covariate because much of the difference between treatments will be due to individual differences, not to the treatments themselves.

3. Minimize the uncontrolled error variance present in every study. Error variance is a normal part of every study, but it can increase if the integrity of the treatments are not maintained. If the experimenter or survey interviewer does not meticulously follow a script, different participants may receive more or less information and thus respond differently, thereby increasing the error variance.

The final element of the t-test formula in Figure 2 is the sample size (n). To increase the value of t, the goal is to increase the sample size. This is perhaps the simplest and most straightforward element of the research design.

Getting the Most from Your Data Collection Effort

As we discussed in the theory section, the most efficient way to conduct research is to think about and design several studies at the same time. This approach enables you to reuse key theoretical concepts, research design elements, and research materials. With really good planning, you can conduct the data collection for several studies at the same time, thus reducing the time spent in data collection.

In the extreme case, you can actually reuse data that you collect in more than one study. However, it is important that any reuse of data is expressly disclosed to the editor when you submit a paper, so that it is clearly understood what is being done. It is also important that the data being reused examine a different research question from any prior manuscript. Republishing data from the same study to test essentially the same research question is not ethical. It is ethical (with disclosure) when it is used to test different research questions.

For example, by examining a series of studies from our prior research we can demonstrate how you might reuse or share data across studies. In this data collection effort, we outlined two separate studies, each addressing a unique research question (Figure 3). The focus of Study 1 was to examine the relationship between group communication environment -- face-to-face versus computer-mediated -- and structured conflict method -- devil's advocacy, dialectical inquiry, or baseline (Valacich and Schwenk, 1995a). The focus of Study 2 was to examine the evaluative tone of the devil's advocate -- objective versus carping -- and group communication environment -- nominal (individuals working alone), face-to-face, or computer-mediated (Valacich and Schwenk, 1995b).

Figure 3 shows the cells used for Study 1 and 2 appropriately labeled and outlined. Notice that two overlapping cells are shared across the two studies, but address unique research questions. In all papers, a description of the program of research and the declaration, with citations, that a subset of the data was previously used in an unrelated analysis was included as a part of the original manuscript sent to the editor.

	Individual	Face-to-Face Group	CMC Group
Baseline		1	1
Dilectical Inquiry		1	1
Objective Devil's Adv.	2	1, 2	1, 2
Carping Devil's Adv.	2	2	2

Figure 3. Overlapping Experimental Designs

EXECUTING THE STUDY

Pilot Testing

One of the most important steps in successful quantitative research is the pilot test step. In our experience, too many junior researchers see pilot testing as something that was important for their dissertation (because their major professor required it) but something that is less important for later work because they now understand how to do research. Too often, the pilot tests are omitted or overlooked in the rush to "get the data." Unfortunately, the pilot test is too important to overlook or to rush.

Most research design books present the pilot test as a way to test the research materials (e.g., procedures, scripts, questionnaires). This assertion is true; pilot tests are important in this way. They help the researcher ensure that the study will work from a practical standpoint, and provide preliminary data to

ensure the reliability of the dependent measures. If the questionnaire items are not reliable in the pilot test, they need to be reworked – and pilot tested again! – before actual data collection begins.

However, the most important reason for the pilot test is not to test the research materials. The most important reason for the pilot test is to provide a preliminary assessment of the theory (like a pre-election poll, if you will). There is no sense in continuing the research if the pilot test data does not provide some indication of support for the theory (unless of course, it disproves the theory and one can develop new theoretical advances). Since the pilot tests are so small relative to the overall experiment, it is usually impossible to conduct a statistical test, but it is possible to get some sense of the directions of the means. At a minimum the means should be at least in the same directions as the theory argues. If they are not, you need to revise the research design – and perhaps the theory – and do another pilot test. Once again, as long as you follow the rules of ethical research design, it is impossible to "prove" a bad theory by tweaking the research design over and over again. As an aside, we typically collect two or three data points per cell during pilot testing. From this small sample, one can typically get a good idea if things are working as planned.

In a study of the exchange of information via groupware, for example, Dennis (1996) conducted four pilot tests. In the initial pilot test, the task was too simple in the that it was easily solved and had too little information to produce differences in the amount of information exchanged; all groups, (groupware or verbal) exchanged most of the information and found the correct solution so there was no variation in the dependent measures. The task was made more complex by adding information and making the correct answer more difficult to discover, and the succeeding pilot tests gradually refined the information load.

Collecting the Data

For laboratory experiments, recruiting students to participate in the study is usually straightforward. The simplest approach is to recruit subjects from large introductory business school courses or from required IS courses. We have found that a 2 to 3 percentage point credit for participation is sufficient to

motivate large numbers of students to participate. We always "overbook" every experimental session by recruiting more subjects than we need because our experience suggests that 10-20% of those who sign-up do not show up.

Finding participants for survey research is a critical element that needs to be done with care, because generalizability is the key contribution of survey research. All too often the same "convenience sample" mentality that is acceptable for experimental research is used for surveys. Participants must be drawn from a sampling frame that fits the theory and that exhibits the desired characteristics discussed above.

AN EXAMPLE

Table 2 presents a summary of these principles in the case of Dennis and Kinney (1998). One issue deserves additional elaboration. The original experiment was designed to include a face-to-face cell as a control treatment to be compared against the 2 x 2 media manipulation. However, this cell was removed during the review process at the request of the reviewers and AE. A subsequent article (Dennis, Kinney, and Hung, 1999) that examined gender differences between face-to-face interaction and computer-mediated communication (CMC) was developed that compared the data in this face-to-face cell with the data in the immediate feedback computer conferencing cell. The second paper had a different focus than the first (face-to-face versus CMC and gender differences) but the reuse of data was disclosed to the editor when the paper was submitted. So while the study was not planned with data reuse in mind, it ended up doing so.

Table 2. Designing the Experiment in Dennis and Kinney (1998)

Understanding Strengths and Limitations	• Focus on precision to the abandonment of realism and generalizability, but the tasks were selected to be relevant to the subjects.
Maximizing Publication Potential	**Construct Integrity** • Construct integrity of the independent variables was ensured by using a tightly worded script, the same experimenter and virtually identical rooms. • Construct integrity of the independent variables was tested through the use of questionnaire-based manipulation checks to ensure that the subjects perceived that two tasks differed on equivocality. • Construct integrity of the dependent questionnaire variables was ensured by using validated measures and newly developed measures (each with 6-19 separate questions), pilot testing, and Cronbach's alpha analysis. • Construct integrity of the dependent coded variables (e.g., decision quality) were validated by asking seven acknowledged task experts to perform the task independently, all of whom reached unanimous agreement. **Testing Theoretical Linkages** • Questionnaire items were used to see if the subjects perceived differences in media richness as hypothesized by MRT; that is did multiplicity of cues and immediacy of feedback affect perceived richness. **Designing for Significance** • Treatments were designed to include a clearly noticeable delay in feedback, but a reasonable one. • Subjects were drawn from the same subject pool to be as homogeneous as possible, and a repeated measures design was used so that subjects served as their own control to minimize error variance. • A group-nested-within-treatment term was used as a covariate to further "pull out" error variance and make it systematic. • A fairly large set of 132 subjects were used. **Getting the Most from the Data Collection** • The data from one cell were used with data from a fifth cell (removed from the original study prior to publication) to study gender differences in face-to-face versus computer mediated communication.
Executing the Experiment	**Pilot Testing** • Twelve pilot tests were run because none of the results supported MRT. Each time the experimental materials were revised to provide a more powerful test and each time they provided little support. **Collecting the Data** • Students were recruited from one 300-person section of the introductory MIS course, and received 2% on their course grade for participating. • Two 2-person groups performed the experiment simultaneously but six subjects were recruited for each time period. When more than four students arrived or an odd number of students arrived, the extra students performed the experiment as individuals.

VI. WRITING THE ARTICLE

One critical aspect that usually gets less attention than it deserves is writing the research article. While many books are devoted to experimental design, we have seen very few books and articles that discuss how to present

research. In our opinion, one of the most compelling articles is one that links publishing scientific research to marketing (Peter and Olson, 1983). For better or worse, publishing research is analogous in many respects to marketing a new product, and many of the underlying marketing principles also apply to publishing.

FINDING THE MESSAGE

The single most important aspect about writing up a study for publication is finding the message of the paper. You must be able to explain the paper's message clearly and concisely -- its unique contribution to knowledge. If you can't condense the central message into one or two sentences, you need to rework your ideas until you can.

Conceptually, the message is analogous to the "unique selling proposition" in marketing. That is, why should someone read and act on the information in the paper. It sometimes helps to find metaphors for key ideas to explain key parts of the message. And occasionally, the message will be quite different than what was originally planned.

The message will guide every aspect of the written element of the paper. The message should run throughout the paper and serve as touchstone in deciding what to write and what to omit. In our experience, junior scholars tend to have great difficulty in omitting interesting ideas and conclusions that do not directly contribute to the message. These "red herrings" usually cause problems as they lead the reviewers and editors to expect things that the study cannot deliver. Even though you may have a great new insight, if it is not essential for the message of the paper, it must be omitted. Discarding good ideas is always hard.

PRESENTATION

The way the paper presents the message is important in convincing reviewers and editors that the paper makes a contribution. Good writing cannot cause a poor experiment to be published, but bad writing can easily doom a good study. Good writing starts by doing your homework: read the journal's

instructions for authors and make sure that your paper conforms to the journal's rules.

Like it or not, most IS papers that present quantitative research follow a standard structure (Table 3) that includes no more than 30 pages of text (plus

Table 3. Standard Structure for a Quantitative Article

Element	Typical Length (double-spaced pages)
Introduction • Provides a compelling reason for reading the paper	1
Theory • Explains the relationships among the key constructs and why those relationships exist • Presents specific, testable hypotheses	8-12
Method • Describes the participants • Describes the independent and dependent variables • Describes the procedures	3-5
Results • Presents a brief factual summary of the statistical analyses and draws conclusions as to whether the hypotheses are supported or not • Contains one table of cell means and standard deviations • Contains one table of the results of the statistical analyses, that includes the test statistic (i.e., t, F), p-value and degrees of freedom • Does not interpret the data beyond the hypotheses	1-3
Discussion • Summarizes the results • Explains why the results occurred, clearly distinguishing between interpretation and conjecture • Presents limitations of the experiment • Draws important implications for managers and practitioners • Draws implications for future research	7-12

tables and figures). The structure makes it easy for reviewers, editors, and ultimately readers of the published article to quickly absorb the message and the key elements of the study. While the standard structure provides a basic outline for the paper, there is still a lot of room for interpretation for what really should go in each section. The best approach is to find one or two papers that you really like as exemplars, so you can model your paper after them.

Conducting Research in Information Systems by A.R. Dennis and J.S. Valacich

Most of the items in a table should be straightforward, but the method section may warrant additional discussion. The method section should be very precise and provide sufficient detail so that another researcher would be able to replicate the study. The rationale behind any important design decisions also should be presented. Never sidestep methodological limitations; the reviewers are usually smart enough to see them anyway. If you know there are limitations, address them head on and explain why you made the choices you did.

Many of the guides to good writing can be useful for academic writing as well (e.g., Strunk and White, 1995). One key to academic writing is succinctness. When in doubt, keep it short. When we edit the first papers written by novice scholars who are just starting their careers, we often can make significant improvements by simply removing words; try it sometimes! Two other common mistakes include the element of surprise and an amateur style. Surprise has no place in academic writing; you are not writing a mystery novel. Every key element of the paper should be foreshadowed (either explicitly, which is simpler, or implicitly). An amateur style includes elements such as the frequent use of quotations, overdone straw man arguments to be dismantled by the paper, exaggeration to increase the intended importance of the paper, the use of old references that are superceded by newer research, the citation of unpublished sources such as dissertations, and the criticism of prior research[2].

DEVELOPING IMPLICATIONS

The discussion section plays an important role in a quantitative research study by drawing implications. The discussion section should be rich in explaining the study, and more importantly, in going beyond the data to draw implications and conclusions.

Three sets of implications are important.

1. The results need to be explained and alternative theory-based explanations developed for results that were unexpected.

---

[2] There is no value in criticizing the work of others because no research is perfect. The goal is to improve prior work. Remember, some of these researchers may end up as reviewers of your paper.

2. Implications for future research need to be presented. What are the next studies that need to be done because of these results? What are some unanswered questions that the study raises that could be answered with empirical data?

3. Implications for managers based on the theory (not the data) need to be developed. What actions should users, managers, systems developers, and so on take because of the results of the study? What should they do or not do? What would you include in your next MBA or undergraduate course based on the study?

We are often disappointed by the mechanical and uninsightful implications found in many initial versions of papers. Our challenge to the readers of this article is to think more carefully about the implications. This is your opportunity to be insightful and say what you really think. Speculate, but don't be boring.

TARGETING

Selecting a target journal and shaping a paper for that journal are important, but can be overemphasized. First and foremost, is the issue of fit. Certain journals tend to publish certain types of papers, and while there are usually no hard and fast rules, some journals are more predisposed to certain types of articles than others, although these predispositions can and do change as editors come and go. The best idea is to go through the last three years of a journal to see what types of articles that have published and whether the editor's comments indicate a desire for more (or less) of a certain type of article. If you find articles in similar style and similar in topic, it suggests the paper may be a good fit -- although you are unlikely to find perfectly related articles.

Once you have selected a journal as a target, the next step is to subtly shape the paper for the journal. Often this will simply involve slightly modifying the introduction and the discussion and conclusion at the end of the paper. You should think about who are likely to be the reviewers and editors and make sure that your paper is aware of their recent work, especially if it has appeared in the journal to which you are submitting your paper. Gratuitously citing the work of

potential editors and reviewers will never earn you credit (it will be recognized for what it is), but omitting relevant work can be embarrassing.

CULTIVATION

No research study is done until it is in print. It is constantly changing as the authors receive comments from colleagues, reviewers, and editors that help improve and refine the message of the study and the presentation of the message in the article. The three important steps in the cultivation of a top tier journal article are:

1. Prototyped and test the study with colleagues. You should create opportunities to discuss the theory and research design with several colleagues before you collect the data, for the same reasons that you prototype information systems or conduct marketing focus groups while the system or product is still under development. Prototyping and testing before data collection enables you to catch and fix problems while they are still easy to fix. After data collection, it is virtually impossible to fix fatal mistakes. Once the data is collected, you should again solicit comments from colleagues on the initial drafts of the papers, because they enable you to catch and fix problems before they reach the reviewers and editors. Straub, Ang, and Evaristo (1994) offer a checklist that can be useful for evaluating papers both before the data are collected, and after (whether you are about to submit the paper, or are a reviewer on the paper).

2. The paper should be "test marketed" at conferences and with others doing research in the area. Presenting a paper forces you to condense and focus the message -- which often improves the paper as you discover better ways to explain the message. The conference review process and the process of addressing questions from conference attendees helps refine the message, and improve its presentation. It also helps you gauge the potential for publication in a top tier journal. Successful researchers are typically skilled at finding and fixing potential problems prior to submitting the article to a journal, which greatly simplifies the review process.

3. The paper must be managed through the review process. Most reviewers and editors genuinely want to help you improve the paper so that it can

Conducting Research in Information Systems by A.R. Dennis and J.S. Valacich

be published. They too are authors and understand the review process from both sides (Lee, 1999). You should work with the reviewers and editor to address every issue they raise. Be sure to provide a detailed response to the reviewers and editor(s) that responds to each and every issue they raise. Your response should explain exactly how you changed the paper to address the issue (with a reference to the exact page numbers or sections that were changed) or an explanation of why you disagree. It is not uncommon for the response to reviewers to be as long as the paper itself. Your job is to make it simple for the editor and reviewers to understand the evolution of the paper.

AN EXAMPLE

Table 4 presents a summary of these principles applied to Dennis and Kinney (1998). Two points require additional elaboration. First, the paper was originally submitted to a management journal. Because virtually all prior management research on MRT had been field or survey based, most reviewers were field or survey researchers with little experience with experiments. One reviewer in the initial review process was a diehard field researcher who deplored the use of undergraduate students and recommended rejecting the article solely on that basis. Fortunately, the editor completely discounted this opinion.

Second, one important criticism of the initial version of the paper submitted to the management journal was the lack of appropriate theory (Table 4 shows that only six pages were devoted to theory). This version simply presented MRT as developed by Daft and Lengel (1986); it did not attempt to justify or elaborate on Daft and Lengel's arguments. The reviewers were not satisfied with this treatment, so theoretical arguments to support MRT were developed, based on the editor's and reviewers' comments. This same criticism was leveled at the initial version submitted to Information Systems Research (ISR), so the additional theoretical justification added in response to the management journal's reviews was removed and replaced with new theoretical arguments proposed by ISR's AE and reviewers.

Table 4. Writing the Article in Dennis and Kinney (1998)

Finding the Message	• The planned message was a controlled test supporting the performance proposition of MRT. • The message changed after the statistical analyses failed to support MRT. • The planned comparison of a face-to-face condition was dropped as being a red herring.
Presentation	The versions of journal article submitted was organized as follows in terms of the number of pages in each section (double-spaced): |	M1	M2	M3	ISR1	ISR2	ISR3
---	---	---	---	---	---	---
Introduction	1	1	1.5	2	2	2
Theory	6	11	12	9.5	8.5	8.5
Method	4	6.5	6	7	4.5	4.5
Results	2.5	2.5	2.5	2	1.5	1.5
Discussion	7	7	7.5	8	9	8
Total	22.5	28	29.5	27.5	26.5	26	
Developing Implications	• The paper included an analysis of alternative explanations. • The paper included implications for future research and implications for managers.
Targeting	• We originally targeted the paper at a management journal because much of the previous research had appeared in both IS and management.
Cultivation	• The initial project idea and the research design was discussed with colleagues prior to execution. • The results were discussed with colleagues before the first paper was written. • The paper was first presented at HICSS in 1994, where it won a best paper award. • The paper was submitted to a management journal in 1994 where it went through two revisions before being rejected in 1995. The response to reviewers for the two revisions were 27 pages and 28 pages, respectively. • The paper was submitted to *Information Systems Research* in 1996 where it went through two revisions before being accepted in 1997. The responses to reviewers for the two revisions were 28 pages and 2 pages, respectively.

Note: M1,M2, and M3 refers to versions submitted to the management journal. The paper was
then submitted to ISR (ISR1) which requested two revisions (ISR2, ISR3) before accepting it.

VII. CONCLUSION

In this paper, we tried to summarize what we see as the key elements in conducting quantitative research. Clearly there is much more that can be said. One of the best books about conducting research (not just quantitative research) is the one edited by Cumming and Frost (1985), which provides a set of more than two dozen articles by leading scholars on their experiences in conducting organizational research. We found the chapter by Daft (1985) to be particularly helpful.

Publishing in a top tier journal takes patience, attention to detail, and determination. It is a long and often winding journey from the first glimmer of an idea to a published article. Quite often the paper that emerges from the review process is quite different and much stronger than the paper that entered the process (e.g., Lee, 1999). And, it is not uncommon for an article to be rejected at a top tier journal, only to find its way into publication -- in a different form -- at another journal. While we have had some success in publishing quantitative research, we have also had our share of papers rejected. Fortunately, the game is one in which we only count successes, not batting averages!

We always believed that research should be fun -- if it's not fun, why do it? In that spirit, and with apologies to David Letterman, we will close with our list of the top ten ways to have your paper rejected at a top tier journal (Table 5, next page). You can use table as a final checklist before you send your paper off to review.

ACKNOWLEDGMENTS

Table 5. Top Ten Ways to be Rejected

1. Avoid theory in favor of a summary of prior research.
2. Omit key papers from your literature review.
3. Include many red herrings that are not related to the paper's message.
4. Plagiarize from the reviewers' articles[3]
5. Openly and directly criticize the work of others -- likely the reviewers of your paper.
6. Theorize one set of constructs, but measure a different set.
7. Describe your research design in vague general terms.
8. Fail to recognize the faults and limitations of the study.
9. For experiments, abandon precision in favor of realism or generalizability; for surveys abandon generalizability in favor of realism or precision.
10. Draw conclusions that differ from your statistical results.
11. Submit a paper that is more than 35 pages of text.
12. Include typographical errors and fail to format you paper to the journal's standards
13. Write obscurely and repetitively.
14. Avoid sharing ideas with colleagues before you design, write and submit the paper.
15. Respond to reviewers' comments with a one page summary.
16. Have more than 10 items in your top ten list.

REFERENCES

Babbie, E. (1995) *The Practice of Social Research*, Belmont, CA: Wadsworth Pub. Co.

Baroudi, J. J. and W. J. Orlikowski (1989). "The Problem of Statistical Power in MIS Research," *MIS Quarterly*, (13), pp. 87-106.

---

[3] Unfortunately, we have both served as reviewers on papers that have plagiarized our work. Plagiarism is not just unethical, it is just plain stupid. Even if the paper does slip through the review process, once it is published, the plagiarism is usually obvious to anyone who knows the field. Plagiarism is a capital offence, so do it once, and your career is dead.

Bacharach, S. B. (1989) "Organizational Theories: Some Criteria for Evaluation," *Academy of Management Review*, (14), pp. 496-515.

Bailey, K.D. (1987) *Methods of Social Research*, third edition, New York: The Free Press.

Cumming, L. L., and P. J. Frost, (eds.) (1985) *Publishing in the Organizational Sciences*, Homewood, Il: Irwin.

Cook, T. D. and D. T. Campbell, (1979) *Quasi-Experimentation: Design & Analysis Issues for Field Settings*, New York: The Free Press.

Connolly, T., L.M. Jessup, and J.S. Valacich (1990). "Effects of anonymity and evaluative tone on idea generation in computer-mediated groups," *Management Science*, 36(6), pp. 689-703.

Daft, R.L. (1985) "Why I Recommended Your Manuscript be Rejected and What You can Do About it," in L.L. Cumming and P. J. Frost (eds.) *Publishing in the Organizational Sciences*, Homewood, Il.: Irwin, pp. 193-209.

Daft, R. L., and R. H. Lengel, (1986) "Organizational information requirements, media richness and structural design," *Management Science*, (32), pp. 554-571.

Davis, G.B. (1992) "An Individual and Group Strategy for Research in Information Systems," in R. Galliers (ed.) *Information Systems Research: Issues, Methods and Practical Guidelines*, Oxford, UK: Blackwell, pp. 230-252.

Dennis, A. R. (1996) "Information Exchange and Use in Group Decision Making: You Can Lead a Group to Information But You Can't Make It Think," *MIS Quarterly*, (20), pp. 433-455.

Dennis, A. R., K. M. Hilmer, and N.J. Taylor, (1998) "Information Exchange and Use in GSS and Verbal Group Decision Making: Effects of Minority Influence," *Journal of Management Information Systems*, (14)3, pp. 61-88.

Dennis, A. R. and S.T. Kinney (1998) "Testing Media Richness Theory In The New Media: The Effects of Cues, Feedback, and Task Equivocality," *Information Systems Research*, (9), pp. 256-274.

Dennis, A. R. Kinney, S.T. and Hung Y. C. (1999) "Gender Differences in the Effects of Media Richness," *Small Group Research*, (30), pp. 405-437.

DeSanctis, G.R., & Poole, M.S. (1994). Capturing the Complexity in Advanced Technology Use: Adaptive Structuration Theory. Organization Science, 5(2), 121-147.

Hayne, S. and R. Rice (1997) "Attribution Accuracy When Using Anonymity In Group Support Systems," *International Journal of Human Computer Studies*, 47(3):429-450.

Jarvenpaa, S. L., G. W. Dickson, and G. DeSanctis (1985) "Methodological Issues in Experimental IS Research: Experiencea nd Recommendations," *MIS Quarterly*, (9), pp. 141-156.

Kerlinger, F.D. (1986) *Foundations of Behavioral Research*, New York: Holt, Rinehart and Winston.

Kinney S.T. and Dennis, A. R. (1994) "Re-Evaluating Media Richness: Cues, Feedback, and Task," *27th Hawaii International Conference on System Sciences*, Maui, pp. IV: 21-30.

Kinney, S.T, and Watson, R.T. (1992) "The Effect of Medium and Task on Dyadic Communication," *Proceedings of ICIS*, Dallas, TX, pp. 107-117.

Lee, A.S. (1991). Integrating Positivist and Interpretive Approaches to Organizational Research. Organization Science, 2(4), 342-365.

Lee, A.S. (1999) "The Role of Information Technology in Reviewing and Publishing Manuscripts at MIS Quarterly," *MIS Quarterly*, (23:4), pp. iv-ix.

Martin, J. (1982) "A Garbage Can Model of the Research Process," in J.E. McGrath (ed.) *Judgment Calls in Research*, Beverly Hills: Sage, pp. 17-39

McGrath, J.E. (1982) "Dilemmatics: The Study of Research Choices and Dilemmas," in J.E. McGrath (ed.) *Judgment Calls in Research*, Beverly Hills: Sage, pp. 69-80

Mennecke, B.E., Valacich, J.S., & Wheeler, B.C. (2000). The Effects of Media and Task on User Performance: A Test of the Task-Media Fit Hypothesis. *Group Decision and Negotiation,* (9)6, 507-529

Nunamaker, J. F. Jr., A. R. Dennis, J.S. Valacich, and D. R. Vogel (1991) "Information Technology to Support Negotiating Groups: Generating Options for Mutual Gain," *Management Science*, (37), pp. 1325-1346.

Orlikowski, W.J., & Baroudi, J.J., (1991), Studying Information Technology in Organizations: Research Approaches and Assumptions. Information Systems Research, 2(1), 1-28.

Peter, J. P. and J. C. Olson, (1983) "Is Science Marketing?" *Journal of Marketing*, (47) pp. 111-125.

Schwenk, C. and J.S. Valacich (1994) "Effects Of Devil's Advocacy And Dialectical Inquiry On Individuals Versus Groups," *Organizational Behavior and Human Decision Processes*, 59, 210-222.

Straub, D. W., S. Ang, S., and R. Evaristo, R. (1994) "Normative Standards for IS Research," *Data Base*, (25)1, pp. 21-34.

Strunk, W. and E.B. White (1995) *Elements of Style*, Needham Heights, MA: Allyn & Bacon.

Sutton, R. I. And Staw, B. M. (1995) "What Theory is Not," *Administrative Science Quarterly*, (40), pp. 371-384.

Trauth, E. M. and L. M. Jessup (in press) "Understanding Computer-mediated Discussions*:* Positivist and Interpretive Analyses of Group Support System Use," *MIS Quarterly*.

Valacich, J. S., and Dennis, A. R. (1994) "A Mathematical Model of Electronic Brainstorming," *Journal of Management Information Systems*, (11:1), pp. 59-72.

Valacich, J.S. and C. Schwenk (1995a) "Devil's Advocacy And Dialectical Inquiry Effects On Group Decision Making Using Computer-Mediated Versus Verbal Communication," *Organizational Behavior and Human Decision Processes*, 63(2), 158-173.

Valacich, J.S. and C. Schwenk (1995b) "Structuring Conflict In Verbal Versus Computer-Mediated Group Decision Making," *Decision Sciences*, 26(3), 369-393.

Valacich, J. S., Mennecke, B., E, Wachter, R. M., and Wheeler, B. C., (1994) "Extensions to media richness theory: A test of the task-media fit hypothesis," *Proceedings of the Twenty-seventh Hawaiian International Conference on System Sciences*, Maui, pp. VI: 11-20.

Van de Ven, A.H., (1989) "Nothing is Quite So Practical as a Good Theory," *Academy of Management Review*, (14), pp. 486-489.

Watson, H. J., J. W. Satzinger, and S. K. Singh (1994) "Establishing Viable Programs of Research: The Executive Information System Program at the University of Georgia," DataBase (25)3, pp. 40-51.

Weick, K. E. (1995) "What Theory Is *Not*, Theorizing *Is*," *Administrative Science Quarterly*, (40), pp. 385-390.

Wheeler, B.C., and J.S. Valacich (1996) "Facilitation, GSS, and training as sources of process restrictiveness and guidance for structured group decision making: An empirical assessment," *Information Systems Research*, 7(4), 429-450.

Whetten, D.A. (1989) "What Constitutes a Theoretical Contribution?" *Academy of Management Review*, (14), pp.490-495

ABOUT THE AUTHORS

Alan R. Dennis is Professor of Information Systems and holds the John T. Chambers Chair of Internet Systems, named in honor of John Chambers, CEO of Cisco Systems. His primary research focus is collaboration technologies, but also performs research in Internet and web-based system development and system use. Within the general area of collaboration, his research focuses on idea generation/creativity, information exchange in decisionmaking, and individual cognition. His research appears in *Information Systems Research, MIS Quarterly, Journal of Management Information Systems, Academy of Management Journal, Management Science, Communications of the ACM, Journal of Applied Psychology, and Organizational Behavior and Human Decision Processes*. He is an Associate Editor at MIS Quarterly, the Executive

Editor of , and serves as a Track Chair for Internet Technologies for The Hawaii International Conference on System Sciences.

Joseph S. Valacich is the Marian E. Smith Presidential Endowed Chair and The George and Carolyn Hubman Distinguished Professor in MIS at Washington State University. Prior to coming to Washington State, he held faculty appointments at Indiana University and The University of Arizona. He received his Ph.D. from the University of Arizona in 1989, and his M.B.A. and B.S. in computer science from the University of Montana. He worked as a programmer, systems analyst, and technical product manager prior to his academic career. He has conducted numerous corporate training and executive development programs for organizations, including: AT&T, Dow Chemical, EDS, Exxon, FedEx, General Motors, and Xerox. His research interests include Technology Mediated Collaboration and Distance Education. His past research has appeared in numerous publications including *MIS Quarterly, Information Systems Research, Management Science, Academy of Management Journal, Communications of the ACM, Decision Science, Organizational Behavior and Human Decision Processes, Journal of Applied Psychology,* and *Journal of Management Information Systems.*