A Guide for Novice Researchers on Experimental and Quasi-Experimental Studies in Information Systems Research

Yair Levy and Timothy J. Ellis
Nova Southeastern University
Graduate School of Computer and Information Sciences
Fort Lauderdale, Florida, USA

levyy@nova.edu, ellist@nova.edu

Abstract

The main focus of this informative article is to bring attention to experimental research in the field of information systems, especially for novice researchers such as doctoral students. In the past three decades, information systems research has been heavily focused on theoretical model development and testing using survey-based methodology. However, criticism on such an approach has been prevalent. Experimental research has been used extensively in the 'hard' sciences and has provided a solid foundation for advancement in those fields. Incorporating a greater emphasis on experimental studies in information systems research is a route to similar advancements in that domain. Although this paper presents little new information, it attempts to make the wealth of existing information on experiments and quasi-experiments usable by the novice researcher. As such, we start by defining the term *experiment* and argue for its importance in the context of information systems research. We then discuss three key categories of experimental design: lab-experiments, quasi-experiments, and factorial design experiments. In each of the key experimental categories, we provide examples of common type(s) of design. Within the lab-experiment, we explore pretest-posttest with control group and Solomon four-group designs. In the quasi-experiment, we discuss nonrandomized pretest-posttest control group design, control-group time series design, and multiple baseline design. We examine factorial design with a discussion of the ex-post facto type of experiment. We conclude the paper with discussions about importance of increased use of experimental research in information systems and it's relevancy to practice and advancement of knowledge.

Keywords: experimental research, research design, experimental design, lab-experimental design, quasi-experimental design, experiments in information systems research.

Introduction

Einstein was quoted saying that "No amount of experimentation can ever prove me right; a single experiment can prove me wrong" (Calaprice, 2005, p. 291). Consistent with Einstein's quote, it appears that considerable amount of scientific research has progressed using experiments. Additionally, experimental research has proven to be a powerful tool in expanding the scientific body of

Editor: Eli Cohen

knowledge (BoK) (Konda, Rajurkar, Bishu, Guha, & Parson, 1999). Experiment in the context of scientific research is defined as "research in which variables are manipulated and their effects upon other variables observed"(Campbell & Stanley, 1963, p. 1). Experiment refers to research "in which an experimenter having complete mastery can schedule treatments and measurements for optional statistical efficiency, with complexity of design emerging only from that goal of efficiency" (Campbell & Stanley, 1963, p. 1). Leedy and Ormrod (2010) defined experimental research simply as "a study in which participants are randomly assigned to groups that undergo various researcher-imposed treatments or interviews, followed by observations or measurements to assess the effects of the treatments" (p. 108). The noteworthy key to an experiment is the researcher's complete control over the research that enables him or her to randomize the study participants in order to provide better assessment of the treatments provided.

In reality, however, the majority of research conducted, especially in the context of business and educational settings, presents considerable difficulty for the researcher to have the luxury of complete control over the research and the ability to randomize participants. Additionally, the reality of research is that "many experimental situations occur in which researchers need to use intact groups. This might happen because of the availability of the participants or because the setting prohibits forming artificial groups" (Creswell, 2005, p. 297). However, researchers can still uncover fruitful knowledge from conducting a non-true or quasi-experiments. The key difference between experiments and quasi-experiments is in the inability of the researcher to randomize the participants into the measured groups (Leedy & Ormrod, 2010). Given the less rigid requirement for the quasi-experiment compared to true-experimental research, researchers must be aware that quasi-experiments also bring increased threats to validity that must be addressed or explicitly documented. This paper will discuss both true-experiments and quasi-experiments.

Experimental design has been documented for thousands of years with simple experiments done in order to provide evidence in various physical and natural settings. Some well known experiments during the seventeen century include the development of Newton's Laws (Cohen & Whitman, 1999). The use of experiments over the years increased in various fields of science including physical sciences, life sciences, social sciences, and applied sciences. Experiments have been useful in providing evidences and proofs for countless decisions. For example, currently in the context of medicine, the U.S. Food and Drag Administration (FDA) requires all drug manufacturers to conduct experiments, known as 'clinical trials,' in order to get initial approval before drugs can be sold (U.S. Food and Drug Administration, 2009). Unfortunately, too few experiments have been done in the information systems domain over the past three decades, while those that were conducted appeared to concentrate on GDSS and virtual teams, yet criticized for the use of students as participants (Paul, Seetharaman, Samarah, & Mykytyn, 2004).

The essence of this paper is to provide novice researchers a brief review of existing experimental designs commonly used in an attempt to simplify the use of such methodologies. We will review some common experimental designs. Additionally, although some types of experimental design for both lab experiments and quasi-experiments may be conducted without a control group, the threats to internal and external validity of such experiments are substantial (Campbell & Stanley, 1963; Cook & Campbell, 1979). Therefore, we decided to highlight here only the top types of experimental designs in the categories of lab experiments and quasi-experiments due to their increased control of internal and external validity threats. Additionally, it's important to note that our advocacy here of the list of highlighted experimental designs is not an exhaustive review of all experimental designs. Additional types beyond what is covered here also should be reviewed if the experimental settings don't follow those prescribed here. Consultation with seminal sources such as Campbell and Stanley (1963) as well as Cook and Campbell (1979) can be beneficial in experimental designs not touched upon here.

Common Types of Experimental Design

Experimental design includes four research categories. The first two – the lab experiment, also known as 'true-experiment,' and the quasi-experiment, also known as the "field-experiment" – are well known. Although less commonly used, the factorial design and the ex-post facto design are also legitimate experimental approaches. The following sub-sections will briefly discuss the key differences between these categories of experiments and provide some common types of experimental design for each category.

Lab Experiment

Lab experiment, or 'true-experiment', is a type of experimental design where the researcher has a great leverage and control over the study, mainly in the form of selecting the participants and randomly assigning participants and/or events into two or more study groups. Such randomization is monumental in reducing threats to internal validity by attempting to isolate any variations between the groups that are due to chance and not due to any given treatment performed (Leedy & Ormrod, 2010). Additionally, the sample selected for the study should be as homogeneous as possible in an attempt to provide additional validity for the measured effect of the treatment. For example, in medical lab experiments for drugs, a researcher may use mice that were breed as near-identical siblings. In the case of IT experiments, especially in the context of research involving people, obtaining "identical participants" is somewhat difficult to obtain. However, the researcher may need to find participants that are as similar as possible in their characteristics known to be relevant to the measured treatment. For example, a proposed lab experiment may attempt to measure the impact of media richness on individuals' propensity to shop online. However, computer self-efficacy (CSE) and social economical status (SES) have been known to produce some impact on individuals' propensity to shop online. Consequently, the researcher would attempt to select study participants that have very similar levels of CSE and are within the same SES. By doing so, the researcher ensures that the study participants are as similar as possible on the given known characteristics (CSE & SES, in this example) that are relevant to the measured characteristic (the individuals' propensity to shop online, in this example).

In lab experiments, researchers exercise a near-full control over the experiment including the randomization of the sample into two groups (experimental and control) and performance of the measurement (M) before the treatment (T), after it, or both. There are two common types of experimental designs for lab experiments, adopted from Campbell and Stanley(1963) as well as Leedy and Ormrod (2010). Both the pretest-posttest with control group design and the Solomon four-group design are known for good control over internal and external validity.

The ***pretest-posttest with control group design*** is the most commonly used experimental design due to its recognized strength in controlling threats to internal validity (Campbell & Stanley, 1963). The researcher randomly assigns participants or events to two groups. The first, the experimental group (Group A), is the group to undergo the prescribe treatment (T_x), while the second, the control group (Group B), is the group that receives no treatment at all and serves as the benchmarking point of comparison. In this design, the researcher performs four measures. Two measures are done prior to the treatment, indicated as 'pretest', one for each of the study groups ($M_{A_{t_1}}$ vs. $M_{B_{t_1}}$). The other two measures are done after the treatment, indicated as 'posttest', one for each of the study groups ($M_{A_{t_3}}$ vs. $M_{B_{t_3}}$). Figure 1 shows the graphical notation of the pretest-posttest with control group design. Ideally in the case of pretest-posttest with control group design, in order to show a valid implication of the treatment on the measures assuming no addi-

tional interference, there should be a significant differences between $M_{A_{t3}}$ vs. $M_{B_{t3}}$, $M_{A_{t3}}$ vs. $M_{A_{t1}}$ andno significant difference between $M_{B_{t3}}$ vs. $M_{B_{t1}}$.

Time (t) →

		t_1 Measure	t_2 Treatment	t_3 Measure
Randomly Assigned	Group A (The Experimental Group)	$M_{A_{t1}}$	T_x	$M_{A_{t3}}$
	Group B (The Control Group)	$M_{B_{t1}}$	-No-	$M_{B_{t3}}$
In an ideal case – desired observed differences		No Diff	-	Sig. Diff
In an ideal case – graphical representation		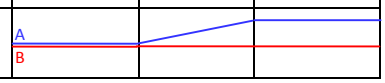		

Figure 1: Pretest-posttest with Control Group Design

An example of pretest-posttest with control group design research may include the testing of the effects of implementing an expert system for IT support group on the time-to-completion of IT troubleshooting tickets. In order to ensure that this is a lab experiment, the researcher likely would conduct the research in a training center, rather than a production IT support shop. In the IT training center, the researcher can decide on the type of individuals that will take part of the research. In doing so, the researcher may select individuals that are demographically the same: age, educational level, experience with technology, culture, etc. The researcher would randomly assign the sample into two groups, experimental and control groups, preferably identical in size and gender distribution. At the start of the research (at time t_1), a measure is made on the time-to-completion of IT troubleshooting tickets for both groups ($M_{A_{t1}}$ vs. $M_{B_{t1}}$), ideally indicating no significant difference between the two. Following (at time t_2), the experimental group will experience the implementation and training of expert system that facilitate assistance with known IT troubleshooting cases, while the control group will not experience such treatment. Then (at time t_3), another measure is made on the time-to-completion of IT troubleshooting tickets for both groups ($M_{A_{t3}}$ vs. $M_{B_{t3}}$), at this point ideally indicating that $M_{A_{t3}}$ is significantly lower then $M_{B_{t3}}$. Given all other variables were under control, the researcher can be confidently conclude that implementation of the expert system for IT support group produced significantly shorter time-to-completion of IT troubleshooting.

The ***Solomon four-group design*** is one of the strongest experimental designs in that it most rigidly controls for threats to both internal and external validity (Campbell & Stanley, 1963). Similar to the pretest-posttest with control group design, in the Solomon four-group design the researcher randomly assigns participants or events to four groups: two experimental groups (Groups A &C) to undergo the prescribe treatment (T_x), and two control groups (Groups B &D), to receive no treatment at all and serve as the benchmarking point of comparison. Unlike the pretest-posttest control group design, however, not all groups are tested prior to the treatment; one of the experimental and one of the control groups is pretested, the other not. The strengths of this experimental design is in its ability to compare not only the differences before the treatment and after the treatment, but also cross reference the comparison with two other groups not measured at the start of the study. The robustness and potential results generalization of the Solomon four-group design results from the fact that the research is also able "to determine how pretesting may affect the final outcome observed" (Leedy & Ormrod, 2010, p. 243). In this design, the researcher per-

forms six measures. Two measures are done prior to the treatment (at time t_1), one for the first experimental group and one for the first control group ($M_{A_{t_1}}$ & $M_{B_{t_1}}$). The treatment (T_x) is provided (at time t_2) to the two experimental groups (A and C). Then (at time t_3), four other measures are done after the treatment, one for each of the study groups ($M_{A_{t_3}}$, $M_{B_{t_3}}$, $M_{C_{t_3}}$, & $M_{D_{t_3}}$).

Figure 2 shows the graphical notation of the Solomon four-group design. Ideally in this case and in order to show a valid implication of the treatment on the measures assuming no additional interference, there should be a significant differences between $M_{A_{t_3}}$ vs. $M_{B_{t_3}}$, $M_{A_{t_3}}$ vs. $M_{A_{t_1}}$, $M_{C_{t_3}}$ vs. $M_{D_{t_3}}$, and $M_{C_{t_3}}$ vs. $M_{B_{t_1}}$, as well as no significant difference between $M_{B_{t_3}}$ vs. $M_{B_{t_1}}$, and $M_{D_{t_3}}$ vs. $M_{B_{t_1}}$.

Time (t) →

		t_1 Measure	t_2 Treatment	t_3 Measure
Randomly Assigned	Group A (Experimental Group 1)	$M_{A_{t_1}}$	T_x	$M_{A_{t_3}}$
	Group B (Control Group 1)	$M_{B_{t_1}}$	-No-	$M_{B_{t_3}}$
	Group C (Experimental Group 2)	-No-	T_x	$M_{C_{t_3}}$
	Group D (Control Group 2)	-No-	-No-	$M_{D_{t_3}}$
In an ideal case – desired observed differences		No Diff	-	Sig. Diff: A|B; C|D No Diff: A|C; B|D
In an ideal case – graphical representation		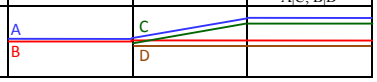		

Figure 2: Solomon Four-Group Design

Quasi-experiment

The quasi-experiment, also known as 'field-experiment' or 'in-situ experiment', is a type of experimental design in which the researcher has limited leverage and control over the selection of study participants. Specifically, in quasi-experiments, the researcher does not have the ability to randomly assign the participants and/or ensure that the sample selected is as homogeneous as desirable. Additionally, in numerous investigations, including those conducted in information systems research, randomization may not be feasible, leaving the researcher with pre-assigned group assignments. Accordingly, the ability to fully control all the study variables and to the implication of the treatment on the study group(s) maybe limited. Never-the-less, quasi-experiments still provide fruitful information for the advancement of research (Leedy & Ormrod, 2010).

An example of a quasi-experiment that does not provide random grouping of participants maybe an investigation of the impact of IT use policy training on employee's IT misuse in an organization. It may very well be that the researcher has no control over which group of employees will receive the training and which group will not as these are based on departments. However, prior research may have indications that employees' computer experience and age have direct implication on employee's IT misuse in an organization (noted as moderator variables or interaction effect). Furthermore, it is very likely that the researcher has very little control over the distribution of the moderator variables (i.e. employees' age and computer experience in this example) between the two groups that may have significant implications for the measure of IT misuse. While

being able to measure IT misuse before the training and after the training, comparing experimental and control groups, even if participants were not randomly distributed, may provide some good indications about the implications of the training, given that the employees' computer experience and age were measured and showed no significant mean differences between the two groups.

Another example of a quasi-experiment is found in the work of Panko (2007) that included measuring the confidence of students in reducing spreadsheet errors and their actual number of spreadsheet errors. His study entailed two quasi-experiments. In one, he compared different types of spreadsheet development (individual development vs. triad development), while in the second, he looked at two groups of students, treatment and control, and measured the confidence level and performance (i.e. reduction in spreadsheet errors) in spreadsheet development before, during, and after spreadsheet training as part of a Management Information Systems (MIS) course. The measures were done at the start of the course, right before the treatment (training/spreadsheet development), after the treatment, and at the end of the course.

The first common type of quasi-experiment is the ***nonrandomized control group pretest-posttest design***, which is similar to the lab experiment's pretest-posttest with control group design but without randomization (Campbell & Stanley, 1963).The experimental group (Group A), is the group to undergo the prescribe treatment (T_x), while the control group (Group B), is the group that receives no treatment at all and serves as the benchmarking point of comparison. In this design, the researcher must be very careful to understand fully the type of treatment provided and what other variables are already known through the literature to be relevant to such treatment and measures. Such knowledge can help the researcher design the appropriate measures even given the limitation of non-randomization. In the example indicated above, for example, prior research may have provided sufficient evidence that employees' age and computer experience has some implications on the measure of IT misuse and/or their understanding of IT use policy. Thus, the researcher in the quasi-experiment should have also measured employees' age and computer experience, comparing such measures to rule out any potential interference of it on the implications of IT use policy training on employees' IT misuse. As in the pure-experimental design, the researcher conducting the nonrandomized control group pretest-posttest study performs four observations. At each of the four observations, there may be multiple variables measured including the key experimental variable (i.e. level of IT misuse in the example above) and the other known and/or suspected moderator variables (i.e. the employees' age and computer experience in the example above). Two observations are done prior to the treatment, indicated as 'pretest', one for each of the study groups ($M_{A_{t_1}}$ vs. $M_{B_{t_1}}$). The other two measures are done after the treatment, indicated as 'posttest', one for each of the study groups ($M_{A_{t_3}}$ vs. $M_{B_{t_3}}$). Figure 3 shows the graphical notation of the nonrandomized control group pretest-posttest design. Similar to the corresponding pure-experiment, in order to refute any differences due to pure chance, there should be a significant differences between the key measured variable $M_{A_{t_3}}$ vs. $M_{B_{t_3}}$, $M_{A_{t_3}}$ vs. $M_{A_{t_1}}$ and no significant difference between $M_{B_{t_3}}$ vs. $M_{B_{t_1}}$. At the same time, any of the controlling variables known and/or suspected should be measured and compared between the two groups. Ideally, the controlling variable(s) across the two groups should show no significant mean differences. The measure of the controlling variable(s) can either be done at t_1, at t_3, or at both times (t_1& t_3) resulting in two or four means comparisons.

		Time (t) → t_1 Measure	t_2 Treatment	t_3 Measure
Non-randomly Assigned	Group A (The Experimental Group)	$M_{A_{t_1}}$	T_x	$M_{A_{t_3}}$
	Group B (The Control Group)	$M_{B_{t_1}}$	-No-	$M_{B_{t_3}}$
In an ideal case – desired observed differences		No Diff	-	Sig. Diff
In an ideal case – graphical representation		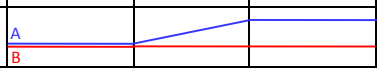		

Figure 3: Nonrandomized Control Group Pretest-Posttest Design

The second common type of quasi-experiment is the ***control group time-series design***, which uses a longitudinal approach (i.e. time-series) of measuring to provide a more reliable measure of the implications of the treatment on the experimental group (Campbell & Stanley, 1963). In this quasi-experimental design there are also two groups: the experimental group (Group A), is the group to undergo the prescribe treatment (T_x), while the second, the control group (Group B), is the group that receives no treatment at all and serves as the benchmarking point of comparison. In this design, the researcher starts by conducting two pretest measures for each group, one at time t_1 ($M_{A_{t_1}}$ & $M_{B_{t_1}}$) and shortly after at time t_2 ($M_{A_{t_2}}$ & $M_{B_{t_2}}$). The treatment (T_x) is provided (at time t_3) to the experimental group (A). Then, the researcher conducts two posttest measures for each group, one at time t_4 ($M_{A_{t_4}}$ & $M_{B_{t_4}}$) and shortly after at time t_5 ($M_{A_{t_5}}$ & $M_{B_{t_5}}$).Figure 4 shows the graphical notation of the control group time-series design. Ideally in this case and in order to show a valid implication of the treatment on the measures assuming no additional interference, there should be significant differences between $M_{A_{t_4}}$ vs. $M_{B_{t_4}}$ and $M_{A_{t_5}}$ vs. $M_{A_{t_5}}$, as well as no significant difference between $M_{A_{t_1}}$ vs. $M_{B_{t_1}}$ and $M_{A_{t_2}}$ vs. $M_{B_{t_2}}$.

		Time (t) → t_1 Measure	t_2 Measure	t_3 Treatment	t_4 Measure	t_5 Measure
Non-randomly Assigned	Group A (The Experimental Group)	$M_{A_{t_1}}$	$M_{A_{t_2}}$	T_x	$M_{A_{t_4}}$	$M_{A_{t_5}}$
	Group B (The Control Group)	$M_{B_{t_1}}$	$M_{B_{t_2}}$	-No-	$M_{B_{t_4}}$	$M_{B_{t_5}}$
In an ideal case – desired observed differences		No Diff	No Diff	-	Sig. Diff	Sig. Diff
In an ideal case – graphical representation		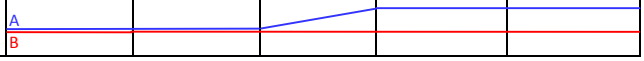				

Figure 4:Control Group Time-Series Design

The third common type of quasi-experiment is the ***multiple baseline design***, which is based on the robust longitudinal approach (i.e. time-series) extended over a longer duration than the pre-test-posttest design (Campbell & Stanley, 1963).In this type of quasi-experimental design, the researcher also has limited control, if any, over the assignment of individuals to the groups. However, the measurement of two groups over time with dual treatment effect may allow for stronger results. The caveat here to the treatments is that such treatment should not be repeated too close in

time allowing the measures to properly reflect the effects of the treatments, while the treatment should be similar in its magnitude. An example of such repeated treatment maybe information security awareness training about corporate policies and differences resulting from employees' password and computer authentication practices. The time between the training sessions (i.e. treatments) can be several months. Similar to other quasi-experimental approaches, there are two groups: the experimental group (Group A), is the group to undergo the prescribed treatments (T_x) twice, while the second, the control group (Group B), is the group that receives the prescribe treatment (T_x) only once and serves as the benchmarking point of comparison. In this type of design, two observations are done prior to the first treatment, one for each of the study groups ($M_{A_{t_1}}$ & $M_{B_{t_1}}$) at the start of the at time t_1. The first treatment (T_{x1}) is provided at time t_2 only to the experimental group (Group A), while a follow up measure is done at time t_3 for both groups ($M_{A_{t_3}}$ & $M_{B_{t_3}}$). Following at time t_4, a second similar treatment (T_{x2}) is provided (such as the information security training example discussed above) both for experimental group and the control group. A final measure is then done at timet_5 ($M_{A_{t_5}}$ & $M_{B_{t_5}}$). Figure 5 shows the graphical notation of the nonrandomized multiple baseline design. Similar to the previous designs, and in order to refute any differences due to pure chance, there should not be a significant difference on the key measured variable(s) $M_{A_{t_1}}$ and $M_{B_{t_1}}$, whereas ideally in this case and in order to show a valid implication of the treatment on the measures assuming no additional interference, there should be significant differences between the measures $M_{A_{t_3}}$ and $M_{B_{t_3}}$, as well as no significant difference between the measures $M_{A_{t_5}}$ vs. $M_{B_{t_5}}$.

Time (t) →

		t_1	t_2	t_3	t_4	t_5
		Measure	Treatment	Measure	Treatment	Measure
Non-randomly Assigned	Group A (The Experimental Group)	$M_{A_{t_1}}$	T_{x1}	$M_{A_{t_3}}$	T_{x2}	$M_{A_{t_5}}$
	Group B (The Control Group)	$M_{B_{t_1}}$	-No-	$M_{B_{t_3}}$	T_{x2}	$M_{B_{t_5}}$
In an ideal case – desired observed differences		No Diff	-	Sig. Diff	-	No Diff
In an ideal case – graphical representation						

Figure 5:Multiple Baseline Design

Factorial Design

Another experimental design is the ***ex-post facto with control group design***, which is the experiment design approach most subject to threats to validity. Never-the-less, there are research studies in which the pre-event or pre-experience measure is either just not feasible or was not performed due to the researcher's inability to predict the need for such a measure. Research to measure the effects of harmful events, where the measure was not evident prior to the event or experience, using the ex-post facto design with control group design can lead to interesting results. Figure 6 shows the graphical notation of the ex-post facto design with control group design.

		Time (t) → t_1 Event/Experience	t_2 Measure
Non-randomly Assigned	Group A (The Event Group)	E_x	$O_{A_{t_2}}$
	Group B (The Control Group)	-No-	$O_{B_{t_2}}$
In an ideal case – desired observed differences		-	Sig. Diff
In an ideal case – graphical representation		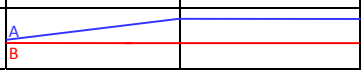	

Figure 6: Ex-post Facto with Control Group Design (Campbell & Stanley, 1963; Leedy & Ormrod, 2010)

An example for an ex-post facto design with control group research is the following: after the 9-11-2001 event, there were attempts to measure the survivability of the IT infrastructure of the U.S. financial companies in New York's Wall Street area. The measure of such survivability prior to this harmful event may not have been performed due to the researcher's inability to predict the need for such a pre-event measure. Comparing such measures of post-event between the IT infrastructure of the U.S. financial companies housed in Wall Street's area (Event Group) and those housed in Chicago (Control Group) may provide considerable interesting findings about the nature of the survivability of the IT infrastructure of the U.S. financial industry due to harmful events.

Assumptions in Experimental Research

Characteristics and Typical Assumptions

Every research conducted includes a set of assumptions (Ellis & Levy, 2009). For each given type of experimental design, the level of each of the assumptions noted below will vary. Therefore, it is the prime responsibly of the researchers to first understand the assumptions under which their research is conducted, and then be candid about such assumptions and properly disclose these assumptions so other researchers wishing to replicate their work can fully understand the differences and/or similarities. The following are known assumptions for experimental research adapted from Leedy and Ormrod (2010, p. 284):

- Random assignment and group composition
- Control over extraneous and moderating variables
- Manipulation of the treatment conditions and consistency of treatments
- Outcome measures
- Group comparison
- Treats to validity

Conclusion

This paper was design as a starting point for novice researchers, especially doctoral students, on few common types of experimental design and their place in IS research. Experiments have shown to be a robust research methodology in mature sciences including applied and behavioral

sciences with substantial advancements in the BoK resulting from both lab-experiments as well as quasi-experiments. Yet, although experimental research has been used rigorously in other fields of research, the use of experiments in IS is somewhat limited. Additionally, with the increase criticism on the relevancy of IS research, especially research conducted in North America, a strategic move towards IS research that is more relevant for industry is highly warranted. We find the following components as critical outcomes of any experimental design conducted in order to be relevant for practice and advancement of knowledge:

- Nature of conclusions drawn in experimental research
- Implications of the experimental research for practice
- Implications of the experimental research for research and recommendations for future experimental research
- Explicit assertion of the experimental research limitations

References

Calaprice, A. (2005). *The new quotable Einstein*. Jerusalem & USA: Hebrew University of Jerusalem & Princeton University Press.

Campbell, D. T., & Stanley, J. C. (1963). *Experimental and quasi-experimental design for research*. Hopewell, NJ: Houghton Mifflin Company.

Cohen, I. B., & Whitman, A. (1999). *Isaac Newton: The principia* (a translation). Berkeley: University of California press.

Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design & analysis issues from field settings*. Boston, MA: Houghton Mifflin Company.

Creswell, J. W. (2005). *Educational research: Planning, conducting, and evaluating quantitative and qualitative research* (2nd ed.). Upper Saddle River, NJ: Pearson.

Ellis, T. J., & Levy, Y. (2009). Towards a guide for novice researchers on research methodology: Review and proposed methods. *Issues in Informing Science and Information Technology, 6*, 323-337. Retrieved from http://iisit.org/Vol6/IISITv6p323-337Ellis663.pdf

Konda, R., Rajurkar, K. P., Bishu, R. R., Guha, A., & Parson, M. (1999). Design of experiments to study and optimize process performance. *International Journal of Quality & Reliability Management, 16*(1), 56-71.

Leedy, P. D., & Ormrod, J. E. (2010). *Practical research: Planning and design* (9th ed.). Upper Saddle River, NJ: Prentice Hall.

Panko, R. R. (2007). Two experiments in reducing overconfidence in spreadsheet development. *Journal of Organizational and End User Computing, 19*(1), 1-23.

Paul, S., Seetharaman, P., Samarah, I., & Mykytyn, P. P. (2004). Impact of heterogeneity and collaborative conflict management style on the performance of synchronous global virtual teams. *Information & Management, 41*(3), 303–321.

U.S. Food and Drug Administration. (2009). *Development and approval process (drugs)*. Retrieved from http://www.fda.gov/Drugs/DevelopmentApprovalProcess/.

Biographies

Dr. Yair Levy is an Associate Professor at the Graduate School of Computer and Information Sciences at Nova Southeastern University. During the mid to late 1990s, he assisted NASA to develop e-learning systems. He earned his Bachelor's degree in Aerospace Engineering from the Technion (Israel Institute of Technology). He received his MBA with MIS concentration and Ph.D. in Management Information Systems from Florida International University. His current research interests include cognitive value of IS, of online learning systems, effectiveness of IS, and cognitive aspects of IS. Dr. Levy is the author of the book "Assessing the Value of e-Learning systems." His research publications appear in the IS journals, conference proceedings, invited book chapters, and encyclopedias. Additionally, he chaired and co-chaired multiple sessions/tracks in recognized conferences. Currently, Dr. Levy is serving as the Editor-in-Chief of the International Journal of Doctoral Studies (IJDS). Additionally, he is serving as an associate editor and as a member of editorial review or advisory board of several refereed journals. Additionally, Dr. Levy has been serving as a referee research reviewer for numerous national and international scientific journals, conference proceedings, as well as MIS and Information Security textbooks. He is also a frequent speaker at national and international meetings on MIS and online learning topics. To find out more about Dr. Levy, please visit his site: http://scis.nova.edu/~levyy/

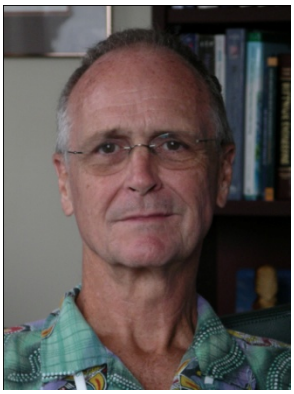

Dr. Timothy Ellis obtained a B.S. degree in History from Bradley University, an M.A. in Rehabilitation Counseling from Southern Illinois University, a C.A.G.S. in Rehabilitation Administration from Northeastern University, and a Ph.D. in Computing Technology in Education from Nova Southeastern University. He joined NSU as Assistant Professor in 1999 and currently teaches computer technology courses at both the Masters and Ph.D. level in the School of Computer and Information Sciences. Prior to joining NSU, he was on the faculty at Fisher College in the Computer Technology department and, prior to that, was a Systems Engineer for Tandy Business Products. His research interests include: multimedia, distance education, and adult learning. He has published in several technical and educational journals including Catalyst, Journal of Instructional Delivery Systems, and Journal of Instructional Multimedia and Hypermedia. His email address is ellist@nova.edu. His main website is located at http://www.scis.nova.edu/~ellist/