Quasi-Experiments

CHARLES S. REICHARDT

Abstract

Quasi-experiments are research designs used to estimate treatment effects when treatments are not assigned at random. Research in quasi-experimentation will advance on four fronts. First, researchers will elaborate the complete array of quasi-experimental comparisons. Second, researchers will refine statistical methods for taking account of initial selection differences. Third, researchers will both improve sensitivity analyses to take account of biases and create empirically based theories of the degree to which biases are removed. And fourth, researchers will assess how well quasi-experiments address the full panoply of complications that arise in practice.

QUASI-EXPERIMENTS

Quasi-experiments are research designs used to estimate the effects of treatments (Shadish, Cook, & Campbell, 2002). Quasi-experiments are widely used because estimating the effects of treatments is a common task and quasi-experiments are easier to implement than other designs, especially in field settings. However, much remains to be known about how quasi-experiments can best be employed to produce high-quality estimates of treatment effects and how to choose the best design and analysis options under different circumstances. Research to answer these questions will focus on (i) the characteristics of the full array of quasi-experimental designs, (ii) the analysis of data from quasi-experiments, (iii) the conditions under which quasi-experiments remove the biasing effects of initial selection differences, and (iv) the ability of different designs to cope with the full range of complications that arise in practice.

For simplicity, only designs that estimate the effect of one treatment compared to a no-treatment or placebo treatment condition will be considered. Generalizing to designs involving more than two treatment conditions is straightforward.

Emerging Trends in the Social and Behavioral Sciences. Edited by Robert Scott and Stephen Kosslyn.
 ISBN 978-1-118-90077-2.

THE ARRAY OF DESIGN OPTIONS FOR ESTIMATING TREATMENT EFFECTS

Estimating the effects of a treatment requires a comparison between what would have happened if the treatment had been implemented and what would have happened if the treatment had not been implemented. Such a comparison can be drawn in a variety of ways. For example, a comparison to estimate the effects of a treatment could be drawn by giving different people different treatments at the same time or by giving the same people different treatments at different times. The effectiveness of the full range of design options has not been well investigated.

Table 1 outlines the fundamental types of randomized and quasi- experimental designs (Reichardt, 2006). The rows distinguish designs where different units of assignment (either participants, times, outcome variables, or settings) receive different treatments. The columns differentiate randomized experiments and two classes of quasi-experiments.

The first row of the table lists designs where participants (e.g., people, animals, classrooms, and cohorts) are the units of assignment. If participants are assigned to different treatments at random, the design is a randomized comparison between participants. Alternatively, participants could be assigned to different treatment conditions based on a cutoff score on a quantitative assignment variable (QAV). Such a design is a quasi-experiment called a *regression-discontinuity* (RD) design (or equivalently a quasi-experimental QAV comparison between participants). In such a design, participants with QAV scores below the cutoff value would be assigned to one treatment condition, while participants with QAV scores above the cutoff value would be assigned to an alternative treatment condition. The outcome variable would be regressed onto the QAV variable in each treatment group. If each of the regression lines were projected to the other side of the cutoff score, the lack of a treatment effect would be evidenced if the two regression lines fell on top of each other. Alternatively, the presence of a treatment effect would be evidenced if one regression line were tilted relative to the other or if one regression line were shifted up or down relative to the other. A third design option would be to assign participants to different treatments neither at random nor according to a QAV. Such a design is a quasi-experiment called a *nonequivalent comparison group* (NEG) design (or equivalently, a quasi-experimental non-QAV comparison between participants).

The second row of Table 1 designates designs where the units of assignment are chronological times. To understand such a design, consider a study to assess whether caffeine causes a person to have headache. At random, the person takes either a caffeine pill or a placebo pill each morning for 100 days and assesses his or her degree of headache pain in the afternoon. The effect of

Table 1
A Typology of Comparisons

Units of Assignment	Assignment to Treatments: Randomized Experiments	Assignment to Treatments: Quasi-Experiments: Quantitative Assignment Variable (QAV)	Assignment to Treatments: Quasi-Experiments: Non-Quantitative Assignment Variable (non-QA)
Participants	Randomized comparison between participants	QAV comparison between participants—the regression-discontinuity (RD) design	Non-QAV comparison between participants—the nonequivalent group design (NEGD)
Times	Randomized comparison between times	QAV comparison between times—the interrupted time-series (ITS) design	Non-QAV comparison between times
Outcome variables	Randomized comparison between outcome variables	QAV comparison between outcome variables	Non-QAV comparison between outcome variables
Settings	Randomized comparison between settings	QAV comparison between settings	Non-QAV comparison between settings

caffeine is then assessed by comparing the results for the days on which the caffeine pills were ingested to the results for the days on which the placebo pills were ingested.

Such a design would be a randomized comparison between times. Alternatively, the person could take the placebo pills for the first 50 days of the study and then take the caffeine pills for the next 50 days of the study (or vice versa). Such a design is a quasi-experiment called an *interrupted time-series* (ITS) design (or a quasi-experimental QAV comparison between times). A third option would be to assign the person to take the caffeine and placebo pills on different days neither at random nor according to a cutoff value along the dimension of time. Such a design is a quasi-experimental non-QAV comparison between times.

Now consider the third row of Table 1 which contains designs where the units of assignment are outcome variables. To understand such designs,

imagine the makers of an educational TV show want to compare two ways of teaching children the letters of the alphabet. The producers of the show divide the letters of the alphabet in half at random and assign one half to be taught using one method of instruction and the other half to be taught using the other method. A large group of children are then exposed to both sets of instructions and the relative effects of the two methods of instruction are assessed by comparing the performances of the children on the two randomly assigned sets of letters. In such a comparison, performances on the different letters are different outcome variables and the design is a randomized comparison between outcome variables. If the letters were assigned to treatment groups based on a cutoff score on a QAV (rather than being assigned to treatment groups at random), then the design would be a quasi-experimental QAV comparison between outcome variables. For example, letters could be ordered based on how frequently they appear in the English language and assigned to treatment conditions according to a cutoff score on that ordering. And, if the letters of the alphabet were assigned to treatment conditions neither at random nor according to a QAV, then the design would be a quasi-experimental non-QAV comparison between outcome variables.

Finally, consider the last row of Table 1 where settings are the units of assignment. To understand such designs, imagine a city that wants to assess the degree to which adding traffic lights to street corners would reduce traffic accidents. If a pool of street corners (to which traffic lights could be added) were available and if traffic lights were added at random to some of the street corners in the pool but not to others, then the design would be a randomized comparison between settings. Alternatively, traffic lights could be assigned to street corners based on a QAV. For example, the street corners could be ordered based on how frequently traffic accidents had occurred during the past 12 months and the street corners with the most accidents could be assigned the traffic lights. Such a design would be a quasi-experimental QAV comparison between settings. Alternatively, if traffic lights were assigned to street corners neither at random nor according to a QAV, then the design would be a quasi-experimental non-QAV comparison.

In practice, research designs are often substantially more complex than the comparisons specified in Table 1. In particular, designs are often combinations of the comparisons presented in Table 1. For example, each of the comparisons in Table 1 could be combined with any of the other comparisons to produce a $4 \times 3 \times 4 \times 3$ set of comparison options (Reichardt, 2009). However, textbooks on quasi-experimentation seldom introduce more than a narrow range of quasi-experimental designs. Indeed, textbooks often introduce only three prototypical quasi-experimental designs: the RD design, the

nonequivalent group design (NEG design), and the ITS design, perhaps along with a few examples of simple design combinations.

Using combinations of quasi-experimental comparisons, rather than a single prototypical design, will often produce the most credible estimates of treatment effects. For example, Yin (2009) describes an evaluation of an innovative middle school program in math and science, where the curriculum was divided into four strands. Schools in the study received instruction in all four strands. A few self-selected schools received innovative instruction in strands 1 and 3, while other self-selected schools received innovative instruction in strands 2 and 4. At the end of the study, the performances of the schools receiving innovative instruction in strands 1 and 3 performed above the average of all the schools on strands 1 and 3 but at the average of all the schools on strands 2 and 4. The results were the opposite for schools that received innovative instruction only in strands 2 and 4. Such a design involved non-QAV comparisons both between participants (i.e., schools) and outcome variables (i.e., strands). Either of these comparisons by itself would have produced results that were not convincing. But when combined, the results were highly credible. Future research will increasingly investigate the effectiveness of designs spanning the full range of options.

ANALYSIS OF DATA FROM QUASI-EXPERIMENTS

In comparisons between participants (comparisons in the first row of Table 1), the participants in the two treatment conditions are not the same. In comparisons between times (comparisons in the second row of Table 1), the chronological times in the two treatment conditions are not the same. In comparisons between outcome variables (comparisons in the third row of Table 1), the outcome variables in the two treatment conditions are not the same. And in comparisons between settings (comparisons in the fourth row of Table 1), the settings in the two treatment conditions are not the same. The initial differences between the units of assignment (either participants, times, outcome variables, or settings) across the treatment conditions are called *initial selection differences*. Differences in the performances of the two treatment conditions could be due to either the effects of the treatments or the effects of selection differences. To estimate the effects of the treatment, the effects of initial selection differences must be removed.

Removing the effects of initial selection differences is relatively easy in randomized experiments. Random assignment guarantees that initial selection differences do not bias the estimate of the treatment effect. In addition, random assignment to treatment conditions makes initial selection differences random which means their effects can be easily bounded within confidence

intervals using simple statistical methods and the bounds can be narrowed simply by increasing the sample sizes.

Without random assignment to treatment conditions, initial selection differences could introduce a bias into the estimate of the treatment effect and it may be difficult to put credible and narrow bounds on the likely size of their effects. Numerous methods and adaptations of methods have been developed to remove the effects of initial selection differences. New statistical methods will continue to be introduced and compared using both real and simulated data. Some of the currently available statistical methods and some of the foreseeable advances in statistical methods are described in the following section.

The Nonequivalent Group (NEG) Design

Statistical methods used to remove the effects of initial selection differences in NEG designs include the analysis of covariance, difference-in-difference estimators, latent variable structural equation modeling, instrumental variable models, Heckman selection models, propensity scores analyses (with different procedures for matching including caliper, kernel density, and nearest neighbor), and doubly robust methods. Future research will attend to three refinements. The first involves measurement error in covariates. Measurement error in covariates can reduce the ability of statistical methods to correct for the effects of initial selection differences. Some statistical methods such as latent variable structural equation models were explicitly designed to take account of measurement error. The development of other methods such as propensity score analyses has largely ignored the problems introduced by measurement error in the covariates. Advances will likely be made to address this oversight in these methods.

Second, in large part individual differences and dose response rates have been given short shrift in estimating treatment effects with NEG designs. Instead, the focus has been on estimating average treatment effects, although differential effects across participants or doses can have important policy implications. The statistical methods that have been developed to analyze data from NEG designs are typically capable of assessing differential effects. However, that capability has often been underutilized. Statistical analysis will more often be exploited to estimate differential effects than they have been in the past.

Third, short shrift has also been given to studying indirect effects which are effects that travel from treatment (X) to outcome (Z) via a specified intermediary variable (Y). Even in a randomized experiment between participants where the assignment of participants to treatments is random, assignment of participants to the intermediary variable (Y) would not be random, so the

comparison used to estimate the effect of the intermediary variable on the outcome would be a quasi-experimental NEG design comparison. Advances will likely be made in the simultaneous analysis of the effects of *X* on *Z*, *X* on *Y*, and *Y* on *Z* in both randomized and NEG designs.

The Regression-Discontinuity (RD) Design

Methods to remove the effects of initial selection differences in RD designs face two significant challenges. The first is assessing the functional form of the regression surface that would appear in the absence of a treatment effect, when the outcome variable is regressed on the QAV. Including polynomial terms in a standard linear regression model or rescaling the outcome or QAV are techniques that have been used to fit curvilinear regression surfaces. More recently, Imbens and Lemieux (2008) have suggested fitting regression surfaces by weighing scores near the cutoff value more heavily than scores farther away. And other statistically sophisticated methods have been developed. Further advances in addressing this problem will be a focus of attention.

The second problem is coping with "fuzzy" assignments to treatment conditions rather than assignments that adhere strictly to the cutoff score on the QAV. Fuzzy assignment is a special case of noncompliance to treatment assignment and has received substantial attention in analyzing data from randomized experiments. Both past and future advances in coping with noncompliance in randomized experiments will likely be applied to the RD design (see also van der Klaauw, 2008). Unfortunately, methods that weigh scores near the cutoff value on the QAV more heavily than scores farther away are at odds with some methods of coping with fuzzy assignment because fuzzy assignment is likely to be most severe near the cutoff value on the QAV.

The Interrupted Time-Series (ITS) Design

The ITS design faces the same challenge as the RD design in estimating the correct functional form of the regression of the outcome variable on the QAV (which in the case of the ITS is chronological time). However, there are also important differences between the ITS design and RD design. The problem of fuzzy assignment appears not to be as widespread in ITS designs as in RD designs, but ITS designs can suffer from the effects of autocorrelation of scores collected over time. A variety of methods have been developed to remove the effects of initial selection differences in the ITS design and, at the same time, account for the effects of autocorrelation. These methods include ARIMA models, multivariate analysis of variance, multi-level models, and

latent variable growth curve models. A potential advance in ITS analysis will be to mirror analysis strategies used in RD designs, including the strategies that weigh more heavily the scores that lie closer to the cutoff value than those that lie farther away. Such mirroring could be especially useful in ITS designs because weighing scores near the cutoff value is not as likely to cause problems due to fuzzy assignment in the ITS design as in the RD design.

The Qualitative Analysis of Data

Another topic that will receive increasing attention in coming years is the implementation of quasi-experimental designs by qualitative researchers. Quasi-experimental methods were developed assuming they would be implemented quantitatively. However, some qualitative researchers assert that the qualitative implementation of quasi-experiments can be superior to their quantitative implementation (Scriven, 2009), which is a conclusion resisted by many quantitative researchers. Nonetheless, a rapprochement between the two camps of researchers has begun and will continue. Qualitative users of quasi-experimental designs must address unique obstacles, such as confirmation biases, as well as show they can cope with all the traditional threats to validity including initial selection differences. And although qualitative researchers often insist their approach to research is based on a different paradigm than the quantitative paradigm, they will discover that the underlying logic of quasi-experimentation (based on drawing imperfect comparisons and ruling out alternative explanations) is shared by both approaches. The trend to use qualitative and quantitative methods together will continue.

UNDER WHAT CONDITIONS AND TO WHAT DEGREE CAN QUASI-EXPERIMENTAL DESIGNS

Remove Bias due to Selection Differences?

All statistical methods devised for estimating treatment effects free from the biasing effects of initial selection differences rest on assumptions. If the assumptions are met, then the statistical methods remove bias due to selection differences. If its assumptions are not met, then a statistical procedure is unlikely to remove bias completely. Unfortunately, the degree to which the necessary assumptions are correct is usually uncertain. If quasi-experiments are to be used to estimate treatment effects, then researchers must know, at least roughly, the degree to which the biasing effects of selection differences

can be removed when the validity of the necessary assumptions is in doubt. Two approaches to obtaining this knowledge are possible.

Sensitivity analysis is one approach to assessing the degree to which quasi-experiments can remove bias due to selection differences. To explicate sensitivity analysis, suppose a statistical procedure perfectly removes the effects of selection differences if the correlation between two variables were precisely zero and, under that assumption, the statistical procedure produces an estimate of a treatment effect with a confidence interval of 14–17 points. Further, suppose the correlation between the two variables is unlikely to be exactly zero but is plausibly believed to lie within a narrow range around zero. Finally, suppose it can be determined that a correlation within the given narrow range around zero would bias the treatment effect estimate, anywhere between −1 and +2 points. Then sensitivity analysis would be said to have shown that the treatment effect, free from the effects of selection differences, is between 12 and 18 points. Implementing sensitivity analyses requires both determining the degree to which the assumptions of the statistical procedures are violated and deriving the effects of those violations on the results of the statistical procedures. The future goal is to derive general ways of accomplishing both tasks. Some advances might be derived from the "uncertainty quantification" of model discrepancies, including Bayesian approaches that incorporate prior distributions of unknown parameters (Brynjardottir & O'Hagan, 2013).

A second approach to determining the degree to which the effects of selection differences can be removed relies on the use of randomized experiments. If a randomized experiment could be implemented free of all biases, then it could be used to estimate the true treatment effect. The results of a quasi-experiment (free of all biases except those due to selection differences) could then be compared to the results from the randomized experiment to assess the degree to which selection biases had not been removed. Studies comparing randomized experiments to NEG designs were attempted in the 1980s but suffered from substantial inadequacies. Improved studies have since compared randomized experiments to both NEG and RD designs (Cook, Shadish, & Wong, 2008). The goal of this research is to derive an empirically based theory of conditions under which quasi-experiments remove bias due to selection differences. So far, some of the tentative conclusions for the NEG design are that its estimates are least biased when the design uses "local" comparison groups that overlap with the treatment group on pretest measures, pretest measures that are operationally identical to the posttest measures, and pretest measures that help determine selection into the treatment groups.

HOW WELL DESIGNS ESTIMATE TREATMENT EFFECTS IN THE FACE OF THE FULL PANOPLY OF COMPLICATIONS THAT ARISE IN PRACTICE

As noted previously, randomized experiments have an advantage compared to quasi-experiments in taking account of initial selection differences. The advantage arises because randomized experiments create initial selection differences that are random and random selection differences can be taken into account with greater credibility than can nonrandom selection differences. However, selection differences that are initially random can become nonrandom in the face of differential attrition and noncompliance to treatment conditions, which often occur when randomized experiments are implemented in field settings. So differential attrition and noncompliance reduce (if not eliminate) the advantages of randomized experiments compared to quasi-experiments in taking account of initial selection differences.

Other sources of bias and complications can arise as well in randomized experiments. For example, biases can arise because of confounds that accompany treatment assignment such as resentful demoralization, John Henry effects, and administrative equalization of treatments (Shadish *et al.*, 2002). Other concerns include the degree to which the estimate of a treatment effect can be generalized beyond a particular research setting because, for example, certain types of people refused to participate in a randomized experiment. In addition, randomized experiments may not be as easy to implement or as economical as quasi-experiments. Hence, it is possible that while randomized experiments are better than quasi-experiments at taking account of initial selection differences, randomized experiments may not be superior to quasi-experiments at estimating treatment effects when faced with the full panoply of complications that arise when designs are implemented in field settings. A focus of research will be on developing an empirically based theory of how randomized experiments compare to quasi-experiments in the face of all likely complications.

The creation of such a theory will help resolve a long-standing debate between qualitatively and quantitatively minded researchers. Some quantitatively minded researchers oversell the benefits of randomized experiments because they focus on the relative advantages that randomized experiments have compared to quasi-experiments in coping with initial selection differences. In contrast, some qualitatively minded researchers oversell the benefits of quasi-experiments because they focus on the relative advantages that quasi-experiments can have compared to randomized experiments in coping with complications other than initial selection differences. An

empirically based theory will make clear the conditions under which different designs are preferable without hyping one over another.

A debate has also arisen about the relative merits of different types of quasi-experiments that parallels the debate about the relative merits of randomized experiments versus quasi-experiments. Many quantitatively minded researchers believe quasi-experiments with quantitative assignment to treatment conditions (comparisons in the second column of Table 1) are generally superior to quasi-experimental designs without quantitative assignment to treatment conditions (comparisons in the third column in Table 1) because of the former's presumed superior ability to take account of initial selection differences. However, many qualitatively minded researchers disagree. Of course, whether quasi-experiments with quantitative assignment to treatment conditions are superior or inferior to quasi-experiments without quantitative assignment to treatment conditions depends on the circumstances. What is needed is an empirically based theory of how different quasi-experiments compare under the typical conditions faced in practical applications.

Currently, such a theory suggests that quasi-experiments with quantitative assignment to treatment conditions are generally better able to control for the effects of initial selection differences than are quasi-experiments without quantitative assignment to treatment effects, but the former will generally be harder to implement and their results will be more difficult to generalize. However, such is only the bare bones of a complete theory. We still have much to learn about how different quasi-experiments compare as well as how different statistical procedures compare when used to analyze data from the same quasi-experiment. For example, it would be useful to compare hierarchical linear modeling approaches with propensity score methods in analyzing data from NEG designs that have several waves of pretest measurements. Similarly, it would be useful to compare propensity score analyses to latent variable structural equation modeling approaches in analyzing data from NEG designs when covariates are measured with error. And it will be important to compare quasi-experiments (as well as randomized experiments) in terms of statistical power and precision, and not just bias. For example, even if the estimate of a treatment effect from an NEG design is biased more by initial selection differences than the estimate from an RD design, that disadvantage might be overshadowed if the NEG design's estimate of the treatment effect were more precise. We also need to know when quasi-experiments (especially the NEG design) are best at assessing individual differences, dose responses effects, and mediating effects.

SUMMARY

I have described four directions for research on quasi-experimentation. First, researchers will investigate the full range and complexity of quasi-experimental comparisons because complex designs that incorporate more than one type of comparison generally produce the most credible results. Second, initial selection differences will always be present in any comparison used to estimate treatment effects and these differences must be addressed if treatment effects are to be estimated. New statistical methods and new adaptations of old methods will be developed to cope with the effects of initial selection differences. And methods developed for use with one type of quasi-experimental design (such as the RD design) will likely cross-fertilize the development of methods for other designs (such as the ITS design). Third, statistical methods can fail to take account of the effects of selection differences if the assumptions underlying the methods are violated. To take account of uncertainty about the validity of assumptions, researchers need to refine sensitivity analyses to take account of biases due to initial selection differences and create empirically based theories of the degree to which biases due to selection differences are removed under different conditions. Fourth, other complications can arise besides initial selection differences. We need an empirically based theory of how well designs and their accompanying statistical analyses function when faced with all the complications that are likely to arise in practice. These four tasks increase in difficulty from the first to the fourth, and progress will likely proceed according to difficulty. However, to the extent we cannot answer the fourth, and hardest, question, we cannot well design studies to estimate treatment effects credibly.

REFERENCES

Brynjardottir, J. & O'Hagan, A. (2013). *Learning about physical parameters: The importance of model discrepancy*. Retrieved from http://www.tonyohagan.co.uk/academic/pub.html

Cook, T. D., Shadish, W. R., & Wong, V. C. (2008). Three conditions under which experiments and observational studies produce comparable causal estimates: New findings from within-study comparisons. *Journal of Policy Analysis and Management*, *27*, 724–750. doi:10.1002/pam.20375

Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, *142*, 615–635.

Reichardt, C. S. (2006). The principle of parallelism in the design of studies to estimate treatment effects. *Psychological Methods*, *11*, 1–18.

Reichardt, C. S. (2009). Quasi-experimental design. In R. E. Millsap & A. Maydeu-Olivares (Eds.), *The SAGE handbook of quantitative methods in psychology* (pp. 46–71). Thousand Oaks, CA: Sage.

Scriven, M. (2009). Demythologizing causation and evidence. In S. I. Donaldson, C. A. Christie & M. M. Mark (Eds.), *What counts as credible evidence in applied research and evaluation practice?* (pp. 134–152). Thousand Oaks, CA: Sage.

Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston, MA: Houghton Mifflin.

van der Klaauw, W. (2008). Regression discontinuity analysis: A survey of recent developments in economics. *LABOUR*, *22*, 219–245.

Yin, R. K. (2009). Student achievement data and findings, as reported in MSPs' annual and evaluative reports. *The Journal of Educational and Policy Studies*, *9*, 139–161.

FURTHER READING

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, *81*, 945–960.

Morgan, S. L., & Winship, C. (2007). *Counterfactuals and causal inference: Methods and principles for social research*. New York, NY: Cambridge University Press.

Reichardt, C. S. (2000). A typology of strategies for ruling out threats to validity. In L. Bickman (Ed.), *Research design: Donald Campbell's legacy* (Vol. 2, pp. 89–115). Thousand Oaks, CA: Sage.

Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometric*, *70*, 41–55.

Rubin, D. B. (2004). Teaching statistical inference for causal effects in experiments and observational studies. *Journal of Educational and Behavioral Statistics*, *29*, 343–367. doi:103102/10769986029003343

Shadish, W. R., & Cook, T. D. (1999). Design rules: More steps towards a complete theory of quasi-experimentation. *Statistical Science*, *14*, 294–300.

West, S. G., Cham, H., & Liu, Y. (2014). Causal inference and generalization in field settings: Experimental and quasi-experimental designs. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social psychology* (2nd ed.). New York, NY: Cambridge University Press.

CHARLES S. REICHARDT SHORT BIOGRAPHY

Charles S. Reichardt is a Professor of Psychology at the University of Denver where he has been since he earned a PhD in 1979. His research focuses on the logic of assessing cause and effect, especially in field settings. His work was awarded the Robert Perloff President's Prize of the Evaluation Research Society and the Jeffrey S. Tanaka Award from the Society of Multivariate Experimental Psychology. He is an elected member of the Society of Multivariate Experimental Psychology and an elected fellow of the American Psychological Society.

RELATED ESSAYS

Social Epigenetics: Incorporating Epigenetic Effects as Social Cause and Consequence *(Sociology)*, Douglas L. Anderton and Kathleen F. Arcaro
To Flop Is Human: Inventing Better Scientific Approaches to Anticipating Failure *(Methods)*, Robert Boruch and Alan Ruby
Repeated Cross-Sections in Survey Data *(Methods)*, Henry E. Brady and Richard Johnston
Ambulatory Assessment: Methods for Studying Everyday Life *(Methods)*, Tamlin S. Conner and Matthias R. Mehl
The Evidence-Based Practice Movement *(Sociology)*, Edward W. Gondolf
Meta-Analysis *(Methods)*, Larry V. Hedges and Martyna Citkowicz
The Use of Geophysical Survey in Archaeology *(Methods)*, Timothy J. Horsley
Network Research Experiments *(Methods)*, Allen L. Linton and Betsy Sinclair
Longitudinal Data Analysis *(Methods)*, Todd D. Little *et al*.
Data Mining *(Methods)*, Gregg R. Murray and Anthony Scime
Remote Sensing with Satellite Technology *(Archaeology)*, Sarah Parcak
Digital Methods for Web Research *(Methods)*, Richard Rogers
Virtual Worlds as Laboratories *(Methods)*, Travis L. Ross *et al*.
Modeling Life Course Structure: The Triple Helix *(Sociology)*, Tom Schuller
Content Analysis *(Methods)*, Steven E. Stemler
Person-Centered Analysis *(Methods)*, Alexander von Eye and Wolfgang Wiedermann
Translational Sociology *(Sociology)*, Elaine Wethington