

5 Reflections on the Contributions of Ward Edwards to Decision Analysis and Behavioral Research

Lawrence D. Phillips and Detlof von Winterfeldt

ABSTRACT

This chapter is based on the writings of Ward Edwards and the recollections of two of his graduate students whom he influenced deeply. Larry Phillips was his student from 1960 to 1966 and Detlof von Winterfeldt was his student from 1970 to 1975. Both continued their interactions with Ward until his death in February, 2005. Larry interviewed Ward in February of 2003 to record his early days and contributions to decision analysis and behavioral research. Video clips from this interview were shown at a special session of the Decision Analysis Society meeting in San Francisco in 2005, and the presentation will be posted on the Decision Analysis Society web site. Detlof met Ward together with Ralph Miles almost weekly in 2003 and 2004, mostly planning new activities, like editing this book, but also conducting interviews and discussing the early days of behavioral decision research and his work at the University of Southern California (USC). A videotaped recording of Detlof's interview with Ward can be obtained from USC. Much of this chapter summarizes these interviews and our personal memories of Ward.

CONTENTS

Ward Edwards: Founder of Behavioral Decision Theory

The Early Days

Bayesian Statistics for Psychologists

Probabilistic Information Processing Systems

Early Tests of the SEU Model

Transitions

USC and the Social Science Research Institute

Retirement

Ward Edwards: Founder of Behavioral Decision Theory

In 1954 and 1961 Ward Edwards published two seminal articles that created behavioral decision research as a new field in psychology (Edwards, 1954, 1961). The topics of this research include how people make decisions and how these decisions can be improved with tools and training. In his 1954 Psychological Review article (Edwards, 1954) he introduced the expected utility model to psychologists and he asked if people actually behave this way, balancing the desirability of an outcome against its chance of occurring, as economists had assumed. That paper identified the issues, but it wasn't until Ward's 1961 Annual Review of Psychology paper (Edwards, 1961) that we see in the title, "Behavioral Decision Theory," the formal beginnings of the new field. In just six years, 139 papers relevant to the discipline had appeared, and subsequent exponential growth prevented any comprehensive tracking of research.

Ward fuelled the growth, but not by proselytizing. He was a listener, working hard at trying to understand colleagues and students as they talked about their research ideas, commenting and criticizing where it was needed, typically with good humor, sometimes expressed in limericks. He was helpful to his students, finding hundreds of thousands of dollars over the years in gifts, research contracts and grants, to fund their studies. He was generous in authorship; if his students had done more work on a paper than he, they were made first author. He traveled extensively, seeking out new sets of people to communicate with, and he told Larry that suffering fools and bastards helps a lot, for it avoids making enemies. He didn't seek honors, finding them pleasant if they arrived, but taking most delight

in being included in the reference lists of other people's papers. His own papers are a delight to read. Eschewing jargon, he communicated complex ideas in simple language, using analogies and metaphors, laced with humor.

The Early Days

Ward was born in Morristown, New Jersey, in 1927. His father was an economist, and as a young man Ward enjoyed overhearing the discussions in his home by economists, who were concerned with real-world issues, a concern that stayed with Ward all his life. After studying psychology at Swarthmore College, he went to Harvard as a graduate student the same year B.F. Skinner arrived, giving Ward a worm's eye view of behaviorism. But, being "a snotty-nosed kid, happy to take a disparate point of view," he found it easy to "reject that as what psychology shouldn't do."¹ Ironically, he married Ruth, Skinner's first Ph.D. student.

At Harvard he studied under Fred Mosteller, who introduced him to the work of von Neumann and Morgenstern and the expected utility model. He was also influenced by S. Smith Stevens, whose work on psychophysics examined the relationship between an objective stimulus and the corresponding sensation experienced by a person. Ward thought it odd that the expected utility model assumed non-linear evaluations of money, but linear evaluations of probabilities. In his thesis (Edwards, 1953), he studied people's preferences for gambles that differed in probabilities and found preferences for some probabilities, especially 50-50 gambles, other things being equal.

In his first job, at Johns Hopkins University, Ward asked for a carrel in the economics library, assembled the literature he hadn't read at Harvard, and eventually produced his landmark 1954 paper. In the meantime, he was supposed to be teaching social psychology, a topic of little interest to him, and his less-than-diligent attention to those duties led Clifford Morgan, then the head of the psychology department, to fire him. His direct supervisor,

Walter ‘Tex’ Garner, then on sabbatical, introduced Ward to Arthur W. Melton, who was head of research at the Air Force Personnel and Training Research Center (AFPTRC). Art found a job for Ward in Denver, where he created the Intellectual Functions section, which was expected to be relevant. He discovered a huge field of unsolved problems to which decision theory is relevant, and he later considered that the most fortunate thing that ever happened to him, for it brought him into direct contact with real-world problems. He soon found his *métier*, exploring the intellectual side of the problem being faced, as he had heard his father’s colleagues do. He also played a lot of poker, a game that was to assume increasing relevance in his career.

A personal conflict with the head of the AFPTRC led Melton to leave for the University of Michigan, with an agreement he could take one person with him. He chose Ward, who was not given a faculty appointment, but could teach one course. Art also brought contacts, funds and contracts to a Department that was growing to be the largest in the world. Ward and Ruth found an old farmhouse to live in at Ann Arbor, with their daughter, Tara, newly-arrived son, Page, and two dachshunds, one named Willy, after Wilhelm Wundt, the founder of experimental psychology. It was the beginning of an era fondly remembered by his graduate students, for the heady evening seminars at the farmhouse, exploring in great depth the latest developments, requiring an original discussion paper by a graduate student at each meeting. Or the dinners at the farmhouse, at which his students savored Ruth’s excellent, if often exotic, cooking, with the early arrivals required to light dozens of candles placed on every horizontal surface in living and dining rooms.

In their offices and ‘back at the lab,’ his students had the pick of over 200 graduate students with whom to interact, providing a stimulating intellectual atmosphere supported by a busy social life. These included Amos Tversky, Paul Slovic, Sarah Lichtenstein, and many others who carved out distinguished careers. But it was partly an uneasy time for Ward, as he

had not yet secured a tenured post, and his occasional colorful and forthright behavior led to rumors that he might not achieve it. Paul, Sarah and Larry organized a letter of support, and though we do not know what effect the letter had on the final decision, tenure was granted. As the reputation of the then Engineering Psychology Laboratory grew, it attracted post-doctoral students Lee Roy Beach, Cameron Peterson, and Jim Shanteau, among others.

Bayesian Statistics for Psychologists

Ward, too, brought research funds to the University from several sources, allowing him the flexibility to move activities from one to another. “Ideas could be pursued, and were.” He had reviewed Jimmy Savage’s book, *The Foundations of Statistics*, and was so impressed with the Bayesian point of view that he later said that if he could take only one book with him to a desert island, that would be the book. And on his return, he would ask what new book has Jimmy written. Ward helped to lure Savage to Michigan, to the Mathematics Department, as there was then no statistics department, perhaps fortunately, as there was no particular opposition to the Bayesian point of view.

Ward asked one of his seminar students, Harold Lindman, to write a paper on Bayesian statistics, and on its receipt persuaded Jimmy to engage with him to revise and improve the paper. That began a series of weekly meetings between Ward and Jimmy, resulting in what Ward considered the third of his ‘blockbuster’ papers (after the 1954 paper and the 1961 Behavioral Decision Theory review), “Bayesian Statistical Inference for Psychological Research”, another Psychological Review paper (Edwards, Lindman, and Savage, 1963). The paper showed that classical and Bayesian statistical inferences may diverge, with the classical approach too willing to reject a true null hypothesis, that the classical and Bayesian views are sometimes incompatible, and that certain characteristics of vague prior opinion can justify

taking a uniform prior as the starting point—the still-unappreciated Principle of Stable Estimation. More generally, they concluded:

Adoption of the Bayesian outlook should discourage parading statistical procedures, Bayesian or other, as symbols of respectability pretending to give the imprimatur of mathematical logic to the subjective process of empirical inference.

Ward explained to Larry how they worked together:

We became good friends. He would say this isn't quite right, I would ask him how, he would tell me, and I'd try to fix it. He was a good mathematical consultant, he would listen and listen, ask questions; only when he was satisfied he understood the problem would he make any suggestions. Our interactions were built around the assumption that I would do the work and he would have the ideas. It worked out very well. There might have been more collaborations if he hadn't died when he did. He was fun to work with, bright, light-hearted, willing to listen, all the things you would like.

Probabilistic Information Processing Systems

During this period, Ward and his students conducted several experimental studies to determine how well the Bayesian model described human behavior in revising opinion as new information was received (Phillips et al., 1966; Edwards et al., 1966). Wilson “Spike” Tanner, one of the originators of signal detection theory (SDT) as a theory of human perception, was also at Michigan, and Ward knew that SDT had successfully challenged threshold theories, so perhaps it would also describe human inference. It didn't. After Larry had plotted numerous graphs of the posterior probabilities assessed by subjects in a complex task involving four hypotheses and 12 types of data, he pointed out that the subjects had

indeed paid attention to prior probabilities, the question the experiment had been designed to answer. “Yes,” replied Ward, looking at the graphs, “but they don’t seem to do much after that.” And so, conservatism was born. This finding reinforced Ward’s enthusiasm for a system he had proposed in 1962, that a computer could use Bayes’ theorem to put together the pieces of information, with people providing the likelihood inputs, later dubbed a “PIP” system (for Probabilistic Information Processing), the forerunner of what would later be called a simple Bayesian net, with only one level in the hierarchy (Edwards et al., 1968).

Throughout his life, Ward continued to champion this idea, with his last attempt to demonstrate it shown in his Hailfinder project, a Bayesian net designed to predict severe weather conditions in eastern Colorado (Abramson et al., 1996). He saw the relevance of that early research to the design of Bayesian nets, and to the task of eliciting likelihoods from experts. He recalled a visit to the North American Aerospace Defense Command (NORAD) at Colorado Springs, before the control centre was moved inside a mountain. He was amazed at the complexity of the information gathering, from radars here and abroad, from ships at sea, and many other sources, all displayed on a huge theatre-screen size transparent map of North America, with a four-light threat level indicator in the upper left (only one lamp lit on the day of our visit!). When Ward asked what was done with all this information, the officer escorting us looked puzzled, so Ward asked what was the output. The officer pointed to a red telephone. Later, he asked me, “Do you think the ratio of input to output information should be like that?” This sense that a better way should be found stayed with him all his life.

That sense motivated a massive study at the University of Michigan, using one of Digital Equipment Corporation’s first computers, a PDP-1, an investigation to determine if the PIP idea was viable. It was. The PIP system, pitted against several other ways of obtaining and combining information, consistently reached firmer conclusions earlier on the basis of the same evidence than the other systems, and this finding was replicated in several

experiments in different laboratories. It soon became obvious that inference structures in the real world are hierarchical; they involve intervening indicators and events between the observable data and the object of the inference, and so a program studying hierarchical inference began in the later 1960s, with many studies reported in a special issue of *Organizational Behavior and Human Performance* in 1973. It was the PIP findings that encouraged Dave Schum, then a graduate student at Ohio State, to begin his productive, life-long study of evidence and inference.

Ultimately, the notion of hierarchical inference was superseded by Bayesian networks and, in a decision making context, influence diagrams. Chapters 10 and 11 of this book cover the most recent developments.

Early Tests of the SEU Model

Throughout his Michigan days, Ward spent many evenings playing poker with Art Melton and other colleagues. Little did he know that one day his telephone would ring inviting him to a meeting of people interested in studying gambling behavior. The sponsor was a multi-millionaire lawyer named Charles B.G. Murphy, who later decided to support Ward's work, to the tune of several hundreds of thousands of dollars. Some of this research was carried out at the Four Queens Casino in Las Vegas, a condition of Murphy's providing financial support to the new owner of the casino. As the approval of the Nevada Gambling Commission was needed, a meeting was arranged for Ward with the head of the commission, Dr. Wayne Pearson, whose PhD from Cornell University, Ward discovered, was on gambling behavior. Pearson, as it turned out, had read all of Ward's published work. He was very helpful to the project then, and later on. With the help of Paul Slovic, Sarah Lichtenstein, Amos Tversky, Dave Krantz, Mark Saltzman, and in particular Barbara Goodman, the project went ahead, with a corner of one room devoted to a task giving, truthfully, "the best odds in Las Vegas."

A professional croupier by the name of John Poticello ran the game, using a roulette wheel and a computer, and a variety of specially-developed gambles, that one way or another offered zero expected-value bets. Ward reported:

I learned some interesting things. I learned that the absolute level of the stakes make less difference than I thought they would; people pay attention to the differences, to the structures of the gambles. I also came to be very clear that my original conclusion from my thesis experiments, that people prefer gambles at some probabilities rather than others, showed up very clearly in the results, like preferences for 50-50 gambles.

He also found that the expected value model describes people's choices very well for two-outcome bets, so that model is descriptive of simple situations, good news for decision analysts who break problems into small pieces, though research is sparse on whether this also works for more complex problems. Looking back on this research, Ward told Larry:

It's impossible for me to say how much of the thinking back at Michigan was influenced by this research, but there was a lot of influence for the people on the project, the custodians of the ideas. It's one of those situations in which the paths of influence are complex and not easily traced, but there nevertheless.

In the mid-1960s, Ward and Art Melton joined forces with Paul Fitts, considered by many to be the "father of engineering psychology," and Bill Hays, author of the thoughtful and accessible *Statistics for Psychologists*, to form the Human Performance Center. The mixture of theoretical and applied work conducted by the Center attracted graduate students and post-docs, including Daniel Kahneman, thereby bringing together the team of Kahneman and Tversky.

By the late 1960s, Ward's marriage ended in divorce, and in 1970 he married Sylvia, a lively Brazilian who tolerated but never much liked the cold winters of Ann Arbor. In the meantime, Paul Fitts died suddenly and unexpectedly, Bill Hays became a dean at another university, and stars like Dick Pew moved on. Ward realized that Art Melton's imminent retirement would leave him in charge of the Center, whose focus had become blurred. Recognizing he could probably not maintain the viability of the Center, Ward moved on to become the Associate Director of the Highway Safety Research Institute in 1971 a position he held for only two years, before moving to USC.

Transitions

Detlof von Winterfeldt joined Ward's Engineering Psychology Laboratory in Michigan in the summer of 1970. At this time Ward was still in the midst of developing and implementing probabilistic information procession systems and he also continued the Las Vegas experiments testing the descriptive validity of the SEU model. He continued the first line of research throughout his academic life, trying to engineer solutions to human judgment and decision making problems, very much in the spirit of engineering psychology.

The second line of descriptive research proved frustrating to him. He discovered the usual descriptive deviations from the SEU model, though no other model did much better (in fact, the SEV model, which he favored throughout his life, did very well). Had he been more intrigued by the deviations from the SEU model, he may well have joined forces with Amos Tversky and Danny Kahnemann to develop a new descriptive theory. Instead, he was more interested in making the SEU model work in practice rather than in discovering its descriptive violations. So he gave up on this line of research altogether.

Two things happened around 1970. On the probability side, Tversky and Kahnemann conducted and published their first experiments on probability biases and heuristics (for a

summary, see Kahneman, Slovic, and Tversky, 1982). On the utility side Keeney and Raiffa developed multiattribute utility theory (Keeney, 1968, 1971; Keeney and Raiffa, 1976). Ward was enthusiastic about the ideas inherent in multiattribute utility theory. However, he thought that the Keeney and Raiffa version was too difficult to use in practice. Always the engineering psychologist, he wanted to create a simple version of this method – which later came to be known as SMART – the Simple Multiattribute Rating Technique. This method is still around, though it went through some metamorphoses.

As much as Ward liked the multiattribute utility research, he had fundamental issues with the research on cognitive heuristics and biases in probability and utility judgments:

1. He deeply believed that people, using appropriate tools, could excel in cognitive tasks. He thought that the heuristics and biases work mischaracterized people's abilities. He was especially incensed by an ill chosen characterization of “man as a cognitive cripple” in one of his colleagues' papers.
2. He was concerned that the cognitive illusions literature would be used to argue against the use of decision analysis (in fact, in a review of one decision analysis application, a reviewer wrote that “Tversky and Kahneman have shown that people can't make these sort of probability judgments,” and used this as an argument to reject the paper).

Ward struggled with the heuristics and biases research for many years and he never made piece with this research, which to the present continues to focus on describing behavior rather than improving capability.

USC and the Social Science Research Institute

Ward always liked problem solving and making a difference with applied research. He also liked to manage research. In 1973 he was presented with a unique opportunity to create the

Social Science Research Institute (SSRI) at the University of Southern California, an interdisciplinary institute spanning several schools, with a commitment to bring social science research to bear on important societal problems. Support came from Zohrab Kaprelian, then USC's Provost and Dean of the Engineering School, who Ward found to be a highly successful combatant, and with whom he "hit it off beautifully." Ward was given a budget and a brief to build an organization, which he did, with the help of several very good people, managing rapid growth from 1973 to the mid-eighties. Two of his Michigan students, David Seaver and Detlof von Winterfeldt, joined him as research assistants at USC. The initial research of this small group at SSRI consisted of tests of multiattribute utility and probability assessment methods. The groups also conducted some studies attempting – with mixed success – to validate multiattribute utility and expected utility models in experimental settings.

SSRI grew quickly, partly due to Ward's management style and philosophy. He often said that his job was to find the brightest researchers, support them well, and then get out of their way. Excellent researchers like Robert Kalaba (mathematics), Malcom Klein (criminal justice) and Sarnoff Mednick (genetic research) joined SSRI, because they shared Ward's approach to applied research and because of the supportive environment the institute provided. Another aspect of Ward's management style was that he liked to work on a one-on-one basis avoiding committees or large group meetings. During the ten years Detlof worked at SSRI, he only remembers two large meetings. The first came in the early days of getting to know the existing faculty and staff and the last occurred, sadly, during the demise of the Institute.

Ward's research in the seventies and eighties continued to focus on developing decision analysis tools, now with a major emphasis on using multiattribute utility theory. He developed several versions of SMART and applied them to social problems (Edwards, 1971,

1977). One quite remarkable application was an evaluation of school desegregation plans for the Los Angeles Unified School District (LAUSD), which was under a court order that required it to develop and evaluate alternative plans (Edwards, 1980). This application of SMART involved 144 evaluation criteria – a feat that he wisely never repeated.

Detlof rejoined Ward in 1978 after a three year appointment at the International Institute of Applied Systems Analysis in Austria. The result of this collaboration were two major accomplishments: The publication of “Decision Analysis and Behavioral Research” (von Winterfeldt and Edwards, 1986) and the development of a more formal approach to multiattribute utility analysis in situations involving multiple stakeholders (Edwards and von Winterfeldt, 1987).

The book *Decision Analysis and Behavioral Research* was meant to be an in-depth examination of the behavioral research that was relevant for decision analysis. Ward spent months working on a chapter covering the by then famous work of Kahneman, Slovic, Tversky and others on cognitive errors in human probability and utility judgments. He was unable to suppress his feelings about this research in many early drafts. He wanted to get it right and therefore solicited comments and criticisms of the early drafts by the authors he criticized – and he got an earful. The experience was painful for him, but eventually this chapter became one of the most important ones in the book. Titled “Cognitive Illusions” it put the literature on probability and utility biases into the perspective of many other biases in human cognition. It also emphasized the usefulness of tools and argued that these biases can be overcome with proper use of tools.

The work on decision analysis with multiple stakeholders generated several papers. Perhaps the most important ideas of this work are that a common value structure can be created, even when stakeholders violently disagree about the issues at hand; that conflicts are often about specific value tradeoffs or facts; that conflicts about values can be expressed as

different weights; and that conflicts about facts can be modeled by using judgments from different experts. Most importantly perhaps was the finding that decision analysis can be useful to help multiple stakeholders understand what they agree and disagree about, focus on the things that they disagree about and explore options that are better for everyone involved.

A change in the University Administration caused Zohrab Kaprelian to resign in the mid eighties, and money from the US Government's Department of Justice vanished, resulting in a 50% downsizing over two years for SSRI. In addition, SSRI, which was originally designed to be a self-standing unit reporting to the provost, was moved into the School of Letters, Arts, and Sciences and Ward had to report to the dean of that school. Aside from the financial implications, this also restricted Ward's ability to hire faculty from other schools and reduced his ability to develop interdisciplinary proposals focused on solving urgent social problems. Ward managed, not without pain, in keeping the organization afloat, and it continued to exist until 2002, though in Ward's view it never recovered from the loss of Zohrab.

Retirement

After his retirement in July 1995, Ward remained active in many projects, though as he had been suffering from Parkinson's disease for many years, his energy was increasingly limited. For the next eight years Ward continued a tradition he started at Michigan in 1962: the annual Bayesian Conference. Researchers and decision analysts from around the world gathered to present their latest ideas. Not, as Ward insisted, their hackneyed old papers, but new thoughts, research, ideas, anything relevant to the theory and practice of decision making. You sent in a brief summary of what you wanted to say, and Ward turned it into a humorous title for the agenda. You talked, participants discussed, and we waited for Ward's wise words, typically encouraging, looking beneath the surface for the intellectual issues, usually finding

things you hadn't thought of, but never embarrassing you in front of colleagues. It was all good fun, and intellectually stimulating.

Two weeks before the 35th Bayesian Conference, in 1997, planned as a Festschrift honoring Ward, Sylvia suddenly died. Although Ward was heartbroken, he decided she would have wanted the celebration to go ahead, and it did. Jim Shanteau, Barbara Mellers and Dave Schum (1998) edited the Festschrift volume, *Decision Science and Technology: Reflections on the Contributions of Ward Edwards*, which delighted Ward for the quality of contributions from his world-wide colleagues. Within a year he married Sandy, whose loving care enabled him to continue through to the 41st Conference, in 2003, which was the last. Ward's energy was draining from him, soon confining him to a wheelchair, but his mind was still active, supported by weekly meetings at his house with Detlof and Ralph Miles. The three of them worked on this book, tentatively called *Advances in Decision Analysis*, which they hoped to publish in 2006 by Cambridge University Press. Active to the end, Ward died on February 1st, 2005. We can think of no better tribute to Ward than to finally see "Advances" in print.

References

Abramson, B., Brown, J., Edwards, W., Murphy, A. H., & Winkler, R. L. (1996).

HAILFINDER: A Bayesian system for predicting extreme weather. *International Journal of Forecasting*, 7, 57-78.

Edwards, W. (1954). The theory of decision making. *Psychological Bulletin*, 41, 380-417.

Edwards, W. (1961). Behavioral decision theory. *Annual Review of Psychology*, 12, 473-498.

Edwards, W., Lindman, H., and Savage, J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, 70, 193-242.

Edwards, W. (1971). Social utilities. *Engineering Economist, Summer Symposium Series*, 6, 119-129.

Edwards, W. (1977). How to use multiattribute utility measurement for social decision making. *IEEE Transactions on Systems, Man and Cybernetics, SMC- 7*, 326-340.

Edwards, W., Phillips, L.D., Hays, W.L., & Goodman, B.C. (1968). Probabilistic information processing systems: Design and evaluation. *IEEE Transactions on Systems Science and Cybernetics, SSC-4*, 248-265.

Edwards, W. (1980). Reflections on and criticism of a highly political multiattribute utility analysis. In L. Cobb and M. Thrall (Eds.). *Mathematical frontiers of behavioral and social sciences*. Boulder, CO: Westview Prases, 157-186.

Edwards, W. and von Winterfeldt, D. (1987). Public values in risk debates. *Risk Analysis*, 7, 141-158.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.

Keeney, R.L. (1968). Quasi-Separable Utility Functions, *Naval Research Logistics Quarterly*, 551-565.

Keeney, R.L. (1971). Utility Independence and Preferences for Multiattribute Consequences, *Operations Research*, 875-893.

Keeney, R.L. & Raiffa, H. (1976). *Decisions with multiple objectives*. New York: Wiley.

Phillips, L.D. and W. Edwards (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, 72, 346-55.

Phillips, L.D., Hays, W.L. & W. Edwards (1966). Conservatism in complex probabilistic inference. *IEEE Transactions on Human Factors in Electronics*, HFE-7, 7-18.

Shanteau, J., Mellers, B. & Schum, D. (Eds.) (1998). *Decision research from Bayesian approaches to normative systems: Reflections on the contributions of Ward Edwards*. Norwell, MA: Kluwer.

von Winterfeldt, D. & Edwards, W. (1986). *Decision analysis and behavioral research*. New York: Cambridge University Press.

¹ The statement in quotation marks are direct quotes of Ward's interview with Larry Phillip's in February of 2003.