The Knowledge Engineering Review, Vol. 5: 1, 1990, 5–14

Behavioural decision theory and it's implication for knowledge engineering*

PAUL E. LEHNER AND LEONARD ADELMAN

Information Systems and Systems Engineering, George Mason University, 4400 University Drive, Fairfax, Virginia 22030, USA

Abstract

This paper explores the implications of research results in behavioural decision theory on knowledge engineering. Behavioural decision theory, with its performance (versus process) orientation, can tell us a great deal about the validity of human expert knowledge, and when it should be modelled. A brief history of behavioural decision theory is provided. Implications for knowledge elicitation and representation are discussed. An approach to knowledge engineering is proposed that takes into account these implications.

1 Introduction

Behavioural Decision Theory (BDT) is a branch of psychology that performs research in human judgement and decision making (JDM). Although related to cognitive psychology research in such areas as human problem solving (e.g., Newell & Simon, 1972), memory, natural language understanding (Schank & Abelson, 1977), and the like, it has traditionally been separarated from these areas by two features.

There is an emphasis on tasks that involve quantitive tradeoffs and the integration of subjective judgements (e.g., probability assessment, option selection).

There is a more of a focus on models that characterize JDM performance rather than JDM processes. For instance, BDT researchers often use linear models to characterize behaviour in contexts where it is believed likely that the "internal" JDM process is pattern-based and non-linear.

There is another feature that distinguishes JDM research—it has had little impact on AI. Whereas cognitive psychology research has often led to new approaches to knowledge representation and elicitation, JDM research has consistently been ignored. In retrospect this should not be surprising. Since much of the focus of AI research is specifically to emulate human problem solving (thereby building an "intelligent" system), JDM theories really do have relatively little to contribute.

In recent years, however, there has been a growing interest in developing practical AI-based "expert systems". To be successful, these systems must perform well. Furthermore, many expert systems address problems (medical diagnosis, weather forecasting, military situation assessment, etc.) that are traditionally considered judgement tasks involving some type of probability estimation. With this shift of emphasis, there has re-emerged a growing interest in the relationship between AI and BDT (e.g., Hammond, 1987; von Winterfeldt, 1988; Levi, 1989; Fischhoff, 1989).

In this paper, we explore the relationship between BDT and the emerging discipline of knowledge engineering. The intent of this paper is tutorial. In particular, we will overview the mainstream of research in BDT and will attempt to identify implications of BDT research for the theory and practice of knowledge engineering. Since most of the work in expert systems and knowledge engineering address inference tasks, this paper will focus primarily on human judgement rather than decision making research.

* Preparation of this manuscript was supported in part by NSF Grant No. ISI-8860178 under a subcontract to DAI, Inc.

2 Overview of behavioural decision theory

As a distinct area of research, BDT is about four decades old. Briefly summarized, the "mainstream" history of behavioural decision theory has proceeded as follows.

Early work was inspired by the same foundations as modern statistical decision theory—namely the Von Neuman and Morgenstern (1944) axiomatization of expected utility (EU) theory and Savage's (1952) extension to incorporate subjective probability. What made these and related writings impressive was that they showed that very general theories of decision making, such as Subjective Expected Utility (SEU) theory, could be derived from a small set of "obvious" behavioural principals. For instance, in risky decision problems (i.e., problems involving uncertain outcomes) with known probabilities it seems obvious that a "rational" decision making behaviour should satisfy the following properties:

Transitivity—Interpret $x > y$ [$x = y$] to mean "gamble" x is more [equally] preferred to y. If $x \geqslant y$, $y \geqslant z$, and a least one of the preferences in strict, then $x > z$.

Independence—Let x, y and z represent three gambles and $[xp, y(1-p)]$ represent a combined gamble of playing x with probability p and y with probability $(1-p)$. Suppose $x > y$. A rational agent should also have the preference $[xp, z(1-p)] > [yp, z(1-p)]$, since this is equivalent to a gamble where one must either play z or have the opportunity to choose between x and y.

Continuity—If $x > y$, $y > z$, then there exists probabilities p and q, where $[xp, z(1-p)] > y$ and $y > [xq, z(1-q)]$.

As it turns out, these axioms imply EU theory (Fishburn, 1988). That is, a decision maker that behaves in a way that violates EU theory must also be violating one of these axioms. In BDT jargon, EU theory is a "normative" theory of decision making. It expresses how a rational decision maker would behave. In a similar manner the normative nature of SEU theory can be derived, and along with it, the normative nature of the probability calculus as the only rational calculus for quantifying degrees of belief.

In psychology these results had a profound effect (Edwards, 1954). With the emergence of the probability calculus and SEU theory as normative theories, one could begin to explore the extent to which human JDM behaviour conformed to these normative ideals. Once experimentation began it was quickly discovered that human JDM behaviour consistently "violated" the principles of SEU. One classic sequence of studies in this area focused on "conservatism" (for a review see Edwards, 1968). In experiments where subjects must re-estimate the probability of a hypothesis given new evidence, revised probability judgements were consistently less extreme than the probability estimate derivable from Bayes Rule. Subjective probability updates were conservative when compared to the Bayesian ideal.

Despite the finding that judgemental updates were conservative, and the emergence of some other deviations from normative ideals, the prevailing viewpoint in the 1960s was that except for some calculation deviations, human JDM "approximated" the SEU/Bayesian ideal. This viewpoint was shattered when Kahneman & Tversky (1973) began publishing their experimental results on judgement heuristics and biases. Briefly summarized, Kahneman and Tversky's main result was that human judgement is characterized by a wide variety of "cognitive biases", and furthermore that judgements are often based on simple heuristics that have no apparent relation to the probability calculus and Bayesian updating. In short, any approximation of human JDM behaviour to the SEU/Bayesian ideal is coincidental.

To illustrate, consider the following problem from Tversky & Kahneman (1980).

"A cab was involved in a hit-and-run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data:

(1) 85% of the cabs in the city are Green and 15% are Blue.

(2) A witness identified the cab as a Blue cab. The court tested his ability to identify cabs under

> the appropriate visibility conditions. When presented with a sample of cabs (half of which were Blue and half Green), the witness made correct identifications in 80% of the cases and erred in 20% of the cases.
>
> *Question*: What is the probability that the cab involved in the accident was Blue rather than Green?"

The typical response to this question is about 80%—a result of focusing on the known accuracy rate of the witness. The "correct" answer, which presumes the use of Bayes Theorem to combine both the base rate and the accuracy rate, is 41%. The reason so many people misjudge the probability in this and similar problems is that they often fail to consider base rate information. This is sometimes referred to as the base rate bias.

The publication of the Kahneman and Tversky results started a new trend in JDM research. Researchers began to search, in earnest, for new types of biases (for a review see Hogarth, 1987) and to develop theories of JDM that were not grounded in the SEU ideal (Kahneman and Tversky, 1979; Tversky and Kahneman, 1981). (Discussion of specific biases is reserved for section 3.) This focus on the weaknesses of human judgement reflected a prevailing viewpoint in the JDM community that for tasks involving subjective judgement human performance was, in general, quite poor.

In the last few years, this negative view of human JDM capabilities has mitigated somewhat. There are several apparent causes for this change of view. The first was the discovery of a "citation bias" (Christensen-Szalanski and Beach, 1984). A careful examination of the literature revealed that published experiments where human JDM behaviour exhibit "good" performance are nearly as numerous as experiments where "poor" performance was observed. However, poor performance experiments were cited far more often—giving the impression that poor performance was a routine result. This does not mean that biases do not exist, but rather that they, like most behaviours, are task dependent.

Second, JDM research has traditionally been done in artificial laboratory settings; and a number of researchers have questioned the extent to which they generalize to "real world" decision tasks. Although there is some evidence that they do generalize (e.g., Tolcott et al., 1989) more research needs to be done. Furthermore, the general nature of "real-world" JDM seems to have become a key focus of recent BDT research (e.g. Klein et al., 1989; Mullin, 1989).

Finally, the normative ideals are being reconsidered. As some researchers have noted (e.g. Fishburn, 1988) a theory is labelled "normative" when it is derived from axioms that people agree to call normative. If, however, one can find problems where people insist on a solution that violates the normative theory, then all one has really done is identify inconsistencies in human judgements of what is to be labeled "normative".

In addition to the "mainstream" history cited above, there is a secondary tradition of research in the JDM literature that has also had a substantial impact of the BDT view of human judgement. This tradition focuses on the use of linear models to predict human judgement (see Dawes, 1979 for review; Levi, 1989 for a recent example). This work is based on two rather surprising, but consistent results. First, in probabilistic domains, linear models are often good predictors of human judgements. Even for tasks that appear to be fundamentally pattern-based (e.g. clinical judgement), a regression model of an expert's judgements is usually a good predictor of that expert's future judgements.

Second, linear models of expert judgement usually perform better than the experts from which they were derived. This result even holds for linear models with improper equal-weights. Indeed, as Dawes (1979, p. 573) noted in his historical summary of research in this area

> "... a search of the literature fails to reveal any studies in which clinical judgement has been shown to be superior to statistical prediction when both are based on the same codable input variables ...
>
> Why? Because people—especially experts in a field—are much better at selecting and coding information than they are at integrating it."

One common explanation for this result is, to quote Hogarth (1987, p. 30)

> "... people are motivated to order and make sense of their world, they have tendencies to seek patterns where none exist and to make unjustified causal attributions."

In short, in probabilistic domains people often make pattern-based judgements on the basis of patterns that are inappropriate and reflect a tendency to track random variance. In addition, as noted by Johnson (1988) people also tend to focus on unusual problem features even when they are not very diagnostic. They give disproportionately little weight to mundane evidential features.

These results lead to the development of "bootstrapping" as an approach to building decision models. In bootstrapping, one begins by eliciting from an expert the factors that the expert believes are good predictors of some event. After the factors and their directional impacts have been identified, implement an equal-weights, intuitive-weights or regression-derived linear model using these factors. The resulting model will almost certainly outperform the expert. Of course, this result is task dependent. But for tasks falling into the category of judgement or prediction (i.e. inference) tasks with uncertainty, this result is very robust.

Finally, it should be noted that no one in the linear models tradition is suggesting that linear models are normative, or even that they will perform well. Rather, the suggestion is simply that they usually perform better than human experts.

3 Implications for knowledge engineering

Simply stated, knowledge engineering is the process of eliciting knowledge from human experts and encoding that knowledge for machine use. Usually the desired result of the knowledge engineering process is an "expert system"—a computer-based system that emulates the problem solving and judgement procedures of a human expert.

Described below are some apparent implications of the JDM literature for knowledge engineering. In suggesting these implications one should keep in mind the goals for which an expert system is being built. If the engineering objective is to develop a system which maximizes performance, then these implications should be considered. However, if the goal is specifically to emulate an expert, or capture a person's knowledge, then they have much less relevance.

It is also important to keep in mind that many psychological phenomena are highly context dependent (Hogarth, 1986) and that very little of this research was performed in a knowledge engineering setting. Consequently, although BDT results do have strong implications for knowledge engineering, empirical research should be performed to test these implications—particularly those implications that seem counterintuitive to experienced knowledge engineers.

3.1 Validity of human expertise

The prevailing opinion in the JDM literature seems to be as follows. Expert knowledge is useful, expert judgement often is not. To identify the factors (inputs) relevant to making an assessment, interview an expert. However, expert judgement—the process of weighing these factors to arrive at an overall assessment—should not be trusted. A knowledge engineer should not be hesitant to explore alternative organizations of expert knowledge, or alternative procedures for aggregating individual factor judgements into an aggregate assessment. One approach to doing this will be outlined in section 4 below.

3.2 Judgemental biases

When eliciting judgements from an expert, be aware that those judgements may reflect significant biases. When a rule, judgement procedure, parameter estimate, etc, is elicited consider any biases that may impact that elicitation.

Presented below are some common biases directly relevant to knowledge engineering. Keep in mind that this is a selective and small sample of a great deal of research in this area.

Availability Bias—People often overestimate the probability of an event that is easy to recall or imagine (e.g. Tversky & Kahneman, 1973).

Confirmation Bias—People tend to seek and focus on confirming evidence, with the result that once they've formed a judgement, they tend to ignore or devalue disconfirming evidence (e.g. Wason, 1960; Tolcott et al, 1989).

Frequency Bias—People often judge the strength of predictive relations by focusing on the absolute frequency of events rather than their observed relative frequency. As Einhorn & Hogarth (1978) have shown, information on the nonoccurrence of an event is often unavailable and frequently ignored when available.

Concrete Information—Information that is vivid or based on experience or incidents dominates abstract information, such as summaries or statistical base-rates. According to Nisbett & Ross (1980), concrete and vivid information contributes to the imaginability of the information and, in turn, enhances its impact on inference.

Conservatism—If people are forced to consider base rates, then they often underestimate the predictive value of new information. That is, their revised probability estimates remain too close to the original base rates (Edwards, 1968).

Anchoring and Adjustment—A common strategy for making judgements is to anchor on a specific cue or value and then to adjust that value to account for other elements of the circumstance. Usually the adjustment is insufficient. So once the anchor is set, there is a bias toward that value (Kahneman & Tversky, 1973).

"*Law of Small Number*"—Problems can be framed in such a way that people, including trained statisticians, give undue confidence to conclusions supported by a relatively small amount of data (Tversky & Kahneman, 1971).

Hindsight Bias—This is perhaps the most problematic of all biases (Fischhoff, 1975). After an event occurs people will often claim they predicted the event ("I knew it all along!"), even though prior to the event they were very uncertain.

Fundamental Attribution Error—People tend to attribute success to their own skill and failure to chance or the situation with which they were faced. However, when evaluating the performance of other people, the tendency is to attribute other people's failure to their personality traits, not the situation (see Nisbett & Ross, 1980).

From the knowledge engineering perspective, biases such as these can be quite problematic. At the very least, the knowledge engineer should be aware of their existence and look out for instances where elicited judgements may reflect a bias. For instance, the tendency to ignore base rates is pervasive. A knowledge engineer should therefore be leery of any elicited rules of the form "If observe A then conclude B with certainty 0.8." since this rule is just as likely to be elicited if the base rate of B is 0.05 as it is if the base rate of B is 0.5. Since neither the expert nor the rule reflect base rate information, the rule is probably inappropriate.

The knowledge engineer should also be aware that an expert may be overconfident in his or her judgement capabilities. For instance, the hindsight bias may lead an expert, who is recalling a past scenario, to believe their judgement processes successfully predicted the scenario outcome, when in fact that they were quite uncertain. In addition, the attribution error makes it likely that when their judgement fails, they may not attribute the failure to their own bad judgement.

3.3 Reporting biases

Knowledge elicitation essentially requires that an expert state how he or she would make a judgement in various situations. Consequently, not only may an elicited rule reflect judgemental biases, but it may also reflect biases in self-reporting how he or she would make judgements. For instance, the hindsight bias and attribution error could lead people to be overconfident in the completeness of their own knowledge. Given such overconfidence, one might speculate that experts are likely to report rules reflecting a stronger connection between pre- and post-conditions than is likely to be true of their own judgements. For instance, an expert may suggest that the rule $P(a \mid b\&c\&d) = .8$, even though in a situation where only b, c, and d were known the expert would judge $P(a) = .6$.

The knowledge engineer should also be aware that elicited probability estimates often depend on the nature of the elicitation question and the scale used to measure these estimates (Hogarth, 1975). Identical questions, asked in different ways, can lead to significant reversals in judgement or preference (e.g. Tversky & Kahneman, 1981; Slovic et al., 1988). These results suggest that a knowledge engineer should consider finding alternative ways to elicit the same judgement. A little redundancy may avoid confusion later.

Finally, the knowledge engineer should be aware of the controversy regarding the accuracy of introspection. Until recently, there was no controversy. Cognitive psychologists, particularly BDT researchers, simply viewed introspection as unreliable (e.g. Nisbett & Wilson, 1977). For instance, Hammond (1976, p. 4) states "... verbal representations of introspective reports of covert mental operations are poor representations because of the well-known inaccuracy of introspection; that is, despite the best of intentions, any introspection regarding the basis for a judgement may simply be an incorrect description of the cognitive activity that took place." More recently, however, researchers more directly interested in problem solving processes have taken a somewhat different view. If a person thinks aloud when performing a problem solving activity, then the verbal introspective reports are highly reliable (e.g. Ericson & Simon, 1984). This is a complex issue that extends well beyond the scope of this paper. It is an issue, however, that the knowledge engineer should be alert to.

3.4 Knowledge representation

Psychological research in human problem solving (e.g. Newell & Simon, 1972), memory, and language understanding (Schank & Abelson, 1977) has traditionally had a significant influence on related areas in AI, particularly in the development of knowledge representation schemes. This seems natural, since these research areas focus their efforts on developing models of "internal" psychological processes.

Process tracing and modelling is also prevalent in the JDM literature (see Ford et al., 1989, for examples). However, most JDM models tend to reflect simple mathematical characteristics of human judgement. This is true even when, at first, these models do not seem to be capable of capturing the many rich subtleties of human judgement.

This trend toward simple models is not without foundation. Experimental comparisons of linear and nonlinear, pattern-based models of human judgement have routinely shown the predictive power of the linear model. To quote Dawes & Corrigen (1974, p. 95):

> Linear models have been successfully employed in a variety of contexts. A review of these contexts indicates that they have common structural characteristics: (a) Each input variable has a conditionally monotone relationship with the output; (b) there is error in measurement; and (c) deviations from optimal weighting do not make much practical difference. These characteristics ensure the success of linear models. In fact, linear models are so appropriate in such contexts that *random linear models* (i.e. models whose weights are randomly chosen except for sign) may perform quite well.

More recently, Stewart et al. (1988) compared the performance of human forecasters, an expert system, and simple weighted-sum linear models in a limited information hail forecasting experiment, and ". . . found that the forecasts made by meteorologists were closely approximated by an additive model [and that] the additive model performed as well as the expert system." (abstract).

The JDM results suggest that, for judgement tasks, sophisticated knowledge representation schemes (i.e. complex nonlinear models) will not improve the performance of an expert system, *so long as the representation scheme is being used to capture more human expertise.* Although nonlinear models can easily outperform linear models, the improvement in performance should *not* be attributed to the fact that such models are better models of human expertise. If a knowledge engineer determines from empirical tests that a nonlinear model will perform better than a linear model, or that there are analytical reasons for preferring a nonlinear approach (e.g., probability models with Bayesian updating), then by all means use the superior model. However, if it is discovered that a nonlinear versus linear model is a more accurate model of an expert's judgements, then the knowledge engineer should strongly consider staying with the linear model. Simply note the nonlinearity for future consideration and testing.

4 Knowledge engineering and normative decision theory

The common perception among JDM researchers is that, when compared to normative ideals, human judgement is often flawed. Given this perception it is not surprising that JDM researchers began to explore the development of "decision aids" that would use normative models to assist people in JDM tasks. These aids seemed to be of two general varieties, which we will label "generic" and "domain-specific" decision aids.

The generic aids serve principally to computationally execute normative model equations. For instance, in the Probabilistic Information Processing (PIP) approach an aid would help people apply Bayes rule to update judgements (Edwards et al., 1968). Given a hypothesis of interest (H1) and some evidence (E), the user of a PIP system would enter his or her prior probability, P(H1), and conditional likelihood odds of receiving E given H1 and $\sim$ H1, (i.e. $P(E \mid H1)/P(E \mid \sim H1)$), and the PIP system would calculate the posterior probability for the user. In principal, aids such as PIP could be used in any problem domain that required judgement. (See, however, Edwards et al., in press, for an interesting discussion of fundamental limits on the usefulness of likelihood ratios for calculating posterior odds.)

Domain-specific aids involved the development of complete decision models that could be used repetitively within a domain. For instance, decision aids that implement Bayesian updating on a hierarchical inference network were already common by the early 1970s (e.g. Kelly & Barkley, 1973). In many ways these aids were/are like today's expert systems. The software was generic (like an inference engine) and it could be used to execute any of several models (knowledge bases) which were elicited from extensive sessions with domain experts.

There is however a key difference. Expert systems seem to promote flexibility in knowledge representation—allowing the integration of rules, frames, scripts, schemas, etc. Decision theory-based aids, on the other hand, promote a restrictive approach—requiring that the encoded knowledge conform to a normative framework such as probability theory (Horvitz et al., 1988).

The impact of this difference is felt primarily in knowledge engineering. Flexibility in knowledge representation allows the knowledge engineer to implement a system that reflects human expert's organization of relevant knowledge. Knowledge elicitation is largely unconstrained. Normative frameworks severely restrict flexibility. Once the "functional form" of the model has been established (e.g. coherent probability model) most of the elicitation effort will involve eliciting factors and parameter values.

Although restrictions on knowledge elicitation may seem undesirable, such restrictions do have one strong benefit. Normative frameworks are considered normative because they satisfy axioms of rational JDM. Most of these axioms are testable. For instance, given that an expert P(H), $P(E \mid H)$ and $P(E \mid \sim H)$ one can compare an expert's judgement of $P(H \mid E)$ with the value calculated by Bayes rule. If they disagree, and the expert insists that the calculated value is incorrect, then the knowledge engineer has a basis for modifying the initial probability statements (Adelman et al., 1982). Consequently, using a normative framework, the knowledge engineer can engage in a substantial test and revision process that can lead to many iterations on a knowledge base—long before the system is ever implemented. Without a normative framework, iterative testing of this type is difficult.

Ideally, the result of this iteration process is a JDM model that satisfies the axioms of a normative framework that the expert finds acceptable. Usually (in our experience) such models will look very different from what an unrestricted model would have looked like. This we consider positive. Substantial improvements have been made to a knowledge base without significant effort. The down side of this approach is that one may not iterate to a model acceptable to the expert (after several sessions the expert simply tires of the whole process, says whatever he or she thinks you want to hear, and then proceeds to forget the whole thing).

Our approach (Lehner et al., 1986, Lehner & Adelman, 1988) is something of a compromise between these two extremes. Usually, we start the knowledge engineering process by trying to build a decision analytic model. After an initial model (the result of several iterations), we begin to relax the requirement for consistency with the normative framework—permitting one to add to the knowledge base heuristic rules that are not grounded in a normative framework. For instance, we may begin with a additive MAU model, but later allow for the addition of procedural rules for modifying parameter values. In a Bayesian inference network, one could allow localized non-Bayesian updating for portions of the network.

This approach is based on our opinion that shifts towards a normative framework that are agreeable to an expert will improve the quality of a knowledge base—but that when the expert insists on disagreeing with the normative framework the framework (which ultimately is normative because people agree it is) begins to look less appealing. The result we believe is a knowledge base that looks very different from the experts initial organization of the problem, but is "almost" normative and hopefully acceptable to the expert.

5 Summary

In this paper we have explored the implications of research results in behavioural decision theory on knowledge engineering. Behavioural decision theory, with its performance (versus process) orientation, can tell us a great deal about the validity of human expert knowledge, and when it is best applied. The main result seems to be that for judgement tasks, the primary contribution of expert knowledge is in the identification of relevant factors, not in the specification of how those factors are combined into an aggregate judgement. Regarding the latter, the knowledge engineer should not be overly concerned with developing a knowledge base that does not reflect the expert's inference processes and organization of knowledge. Often, a better model can be developed by deviating from the expert's judgements and conceptual model of the problem. In short, to build a better expert system, the knowledge engineer should be aware of the strengths and weaknesses of human expertise.

References

Adelman, L, Donnell, M, Phelps, R and Patterson, J, 1982, "An iterative Bayesian decision aid: Toward improving the user-aid and user-organization interface" *IEEE Trans. Systems, Man, and Cybernetics* **SMC-12** 733–742.

Christenses-Szalanski, J and Beach, L, 1984, "The citation bias: Fad and fashion in the judgement and decision literature" *American Psychologist* **39** 75–78.

Dawes, R, 1979, "The robust beauty of improper linear models in decision making" *American Psychologist* **34**(7) 571–582.

Dawes, R and Corrigan, B, 1974, "Linear models in decision making" *Psychological Bulletin* **81** 95–106.

Edwards, W, 1954, "The theory of decision making" *Psychological Bulletin* **51** 380–417.

Edwards, W, 1968, "Conservatism in human information processing" in: Kleinmuntz, B, ed., *Formal Representation of Human Judgment* Wiley.

Edwards, W, Phillips, L, Hays, W and Goodman, B, 1968, "Probabilistic information processing systems: Design and evaluation" *IEEE Transactions on Systems Science and Cybernetics* **SCC-4** 248–265.

Edwards, W, Schum, D and Winkler, R, in press, "Murder and (of?) the likelihood principle: A trialogue" *Journal of Behavioral Decision Making*, to appear.
Einhorn, H and Hogarth, R, 1978, "Confidence in judgment: Persistence of the illusion of validity" *Psychological Review* **85** 395–416.
Einhorn, H and Hogarth, R, 1986, "Judging probable cause" *Psychological Bulletin* **99** 3–19.
Ericson, K and Simon, H, 1984, *Protocol Analysis: Verbal Reports as Data* MIT Press.
Fishchhoff, B, 1975, "Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty" *Journal of Experimental Psychology: Human Perception and Performance* **4** 330–344.
Fischhoff, B, 1989, "Eliciting knowledge for analytical representation" *IEEE Transactions on Systems, Man and Cybernetics* **19**(3) 448–461.
Fishburn, P, 1988, "Normative theories of decision making under risk and under uncertainty" in: Bell, D, Raiffa, H and Tversky, A, eds., *Decision Making: Descriptive, Normative, and Prescriptive Interactions* Cambridge University Press.
Ford, J, Schmitt, N, Schechtman, S, Hults, B and Doherty, M, 1989, "Process tracing methods: Contributions, problems and neglected research questions" *Organizational Behavior and Human Decision Making* **43** 75–117.
Hammond, K, 1976, *Facilitation of Interpersonal Learning and Conflict Reduction by One-line Communication* University of Colorado.
Hammond, K, 1971, "Computer graphics as an aid to learning" *Science* **172** 903–908.
Hammond, K, 1987, "Toward a unified approach to the study of expert judgment" in: Mumpower, J, Phillips, L, Renn, O and Uppuluri, V, *Expert Judgment and Expert Systems* Springer-Verlag.
Hogarth, RM, 1975, "Cognitive processes and the assessment of subjective probability distributions" *Journal of the American Statistical Association* **70** 272–289.
Hogarth, R, 1986, "Generalization in decision research: The role of formal models" *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-16**(3) 439–449.
Hogarth, R, 1987, *Judgment and Choice* John Wiley.
Horvitz, E, Brese, J and Henrion, M, 1988, "Decision theory in expert systems and artificial intelligence" *International Journal of Approximation Reasoning* **2** 247–302.
Johnson, E, 1988, "Expertise and decision under uncertainty: Performance and process" in: Chi, M, Glaser, R and Farr, M, eds., *The Nature of Expertise* Lawrence Erlbaum Associates.
Kahneman, D and Tversky, A, 1973, "On the psychology of prediction" *Psychological Review* **80** 237–251.
Kahneman, D and Tversky, A, 1979, "Prospect theory: An analysis of decision under risk" *Econometric* **47** 263–289.
Kelly, C and Barkley, S, 1973, "A generalized Bayesian model for hierarchical inference" *Organizational Behavior and Human Performance* **10** 388–403.
Klein, G, Calderwood, R and MacGregor, D, 1989, "Critical decision method for eliciting knowledge" *IEEE Transactions on Systems, Man and Cybernetics* **19**(3) 462–472.
Lehner, P and Adelman, L, 1988, "Senior battle staff decision aiding: A case study" in: Andriole, S and Hopple, G, eds., *Defense Applications of Artificial Intelligence* Lexington Books.
Lehner, PE, Probus, MR and Donnell, M, 1985, "Building decision aids: Exploiting the synergy between decision analysis and artificial intelligence" *IEEE Transactions on Systems, Man and Cybernetics* **SMC-14**(4), 409–414.
Levi, K, 1989, "Expert systems should be more expert than human experts: Evaluation procedures from human judgment and decision making" *IEEE Transactions on Systems, Man and Cybernetics* **19**(3) 647–657.
Mullin, T, 1989, "Expert's estimation of uncertain quantities and its implications for knowledge acquisition" *IEEE Transactions on Systems, Man and Cybernetics* **19**(3) 616–625.
Newell, A and Simon, H, 1972, *Human Problem Solving* Prentice Hall.
Nisbett, R and Ross, L, 1980, *Human Inference: Strategies and Shortcomings of Social Judgment* Prentice-Hall.
Nisbett, R and Wilson, T, 1977, "Telling more than we can know: verbal reports on mental processes" *Psychological Review* **84** 231–259.
Savage, LJ, 1954, *The Foundation of Statistics* Wiley.
Schank, R and Abelson, R, 1977, *Scripts, Plans, Goals, and Understanding* Erlbaum Associates.
Slovic, P, Fischhoff, B and Lischtenstein, S, 1977, "Behavioral decision theory" *Annual Review of Psychology* **28** 1–39.
Slovic, P, Fischhoff, B and Lichtenstein, S, 1988, "Response mode, framing, and information-processing effects in risk assessment" in: Bell, D, Raiffa, H and Tversky, A, eds., *Decision Making: Descriptive, Normative, and Prescriptive Interactions* Cambridge University Press.
Tolcott, MA, Marvin, FF and Lehner, PE, 1989, "Expert decision making in evolving situations" *IEEE Transactions on Systems, Man and Cybernetics* **19**(3) 606–615.
Tversky, A and Kahneman, D, 1971, "The belief in the 'Law of Small Numbers'" *Psychological Bulletin* **76** 105–110.

Tversky, A and Kahneman, D, 1973, "Availability: A heuristic for judging frequency and probability" *Cognitive Psychology* **5** 207–232.

Tversky, A and Kahneman, D, 1980, "Causal schemas in judgement under uncertainty" in: Fishbein, M, ed., *Progress in Social Psychology* Erlbaum.

Tversky, A and Kahneman, D, 1981, "The framing of decisions and the psychology of choice" *Science* **211** 453–458.

von Neumann, J and Morgenstern, O, 1944, *Theory of Games and Economic Behavior* Princeton University Press.

von Winterfeldt, D, 1988, "Expert systems and behavioral decision research" *Decision Support Systems* **4** 461–471.

Wason, P, 1960, "On the failure to eliminate hypotheses in a conceptual task" *Quarterly Journal of Experimental Psychology* **12** 129–140.