

**Warrants for Prescription: Analytically and Empirically Based
Approaches to Improving Decision Making**

Raanan Lipshitz

Department of Psychology

University of Haifa, Haifa 31905, Israel

e-mail: Raanan@psy.haifa.ac.il

and

Marvin S. Cohen

Cognitive Technologies Inc.,

4200 Lorcom Lane, Arlington, VA 22207

e-mail: mcohen@cog-tech.com

June, 2004

Key words: Naturalistic Decision Making, Descriptive Modeling, Prescriptive Modeling.

The authors wish to thank Rex Brown, Michael Brickner, and Gary Klein for their helpful comments.

Citation:

Lipshitz, R. & Cohen, M. S. (2005). Warrants for prescription: Analytically and empirically based approaches to improving decision making. *Human Factors* 47(1), 121-130.

Abstract

Efforts to improve decision making must appeal to some source of warrant, i.e., specific criteria or models for guiding and evaluating decision making performance. We examine and compare the warrants for two approaches to decision aids, decision training, and consulting: analytically based prescription, which obtains warrant from formal models, and empirically based prescription, which obtains warrant from descriptive models of successful performance. We argue that empirically based warrants can provide a meaningful and valid basis for prescriptive intervention without committing the naturalistic fallacy (i.e., confusing what is with what ought to be) and without the use of formal deduction from first principles. We describe points of divergence as well as convergence in the types of warrant appealed to by naturalistic decision making and decision analysis, letting each approach shed light on the other, and explore the application of empirically-based prescription to Cognitive Engineering.

Introduction

The study of decision making has always been motivated, at least in part, by the desire to improve it. Cognitive engineering is a burgeoning subspecialty within the human factors and ergonomics community that is focused on the design, development, and testing of interventions for that purpose. Of necessity, it requires criteria and models to determine when decision making needs improvement, what kinds of interventions are appropriate, and the interventions' degree of success. Such models and criteria provide *warrants* for prescription.

A large number of present-day decision researchers frame the question of warrants, implicitly or explicitly, in terms of a distinction that first became pervasive in the 1950's (e.g., Savage, 1954) between *descriptive* and *normative* decision making models. Descriptive research is concerned with how people actually make judgments and choices, and it aims at empirical accuracy and (if possible) explanatory adequacy. Normative research tells us how people *ought* to make decisions, by specifying principles and constraints derived from formal or mathematical systems such as deductive logic, Bayesian probability theory, and decision theory (also known as *rational choice theory* or the *subjectively expected utility* (SEU) model). According to researchers in this tradition, the need to improve decision making arises because human decision makers systematically violate normative constraints. Due to limited computational capacity, however, human decision makers cannot directly implement ideal normative models. To bridge the gap between the normative and the descriptive, Bell, Raiffa, and Tversky (1988) introduced the term *prescriptive* to refer to a third category of research, the "engineering side of the pure (normative) theory" (Raiffa, 1994, p. 4; von Winterfeldt, 1999). A prescriptive intervention (such as decision training, decision analytic consulting, or decision aiding) aims to help real people more nearly satisfy the normative ideal within the constraints of their cognitive abilities. Experimental work in behavioral decision making helps identify targets for prescriptive intervention (e.g., Kahneman, Slovic, & Tversky, 1982;

Kahneman & Tversky, 2000; Gilovich, Griffin, & Kahneman, 2002), and experiments in cognitive engineering determine whether decision training, advice, or decision aids successfully *reduce* departures from formal normative standards (e.g., Baron & Brown, 1991; Nisbett, 1993). Because formal normative models provide the warrants for intervention, its content, and the ultimate criterion of success, we call this widely shared paradigm *analytically based prescription*.

Are analytically based normative models the only source of warrants for prescription? Should analytically based models play a dominant role in the design and evaluation of prescriptive practices? One alternative to analytically based prescription is *descriptive* research on the way proficient or expert decision makers accomplish real-world tasks. Such research is conducted by experimental psychologists (Ericsson & Smith, 1991), researchers in judgment and decision making (Shanteau, 1992), and researchers in naturalistic decision making (NDM) (Klein, Orasanu, Calderwood & Zsombok, 1993; Cohen, Freeman, & Wolf, 1996; Lipshitz, Klein, Orasanu, & Salas, 2001). Researchers in naturalistic decision making, in particular, have designed training and decision aids based on discrepancies in the knowledge representations and cognitive strategies used by more and less experienced decision makers (Crandall and Getchell-Reiter, 1993; Cohen & Freeman, 1997; Cohen, Freeman, & Thompson, 1997, 1998; Cohen & Thompson, 2001; Klein, 1997; Pliske, McCloskey and Klein, 2001). Because the warrants for intervention, the contents of prescriptions, and the standards of success are based on empirical research rather than formal models, we call this approach *empirically based prescription*. Our purpose in this paper is to compare analytically-based and empirically-based approaches and to present the broader implications of the latter for prescriptive applications such as system design, decision aiding, and training. These implications have not been discussed systematically in the decision making and cognitive design literatures.

The value and indeed the existence of empirically based prescription is controversial. For example, LeBoeuf & Shafir (2001, p. 374) asked how naturalistic decision making can “presume to offer prescription without a normative framework.” A related question is, How can prescription be based on description without confusing “what is” with “what ought to be” (the naturalistic fallacy)? Bell et al. (1988) reserved exclusive use of the term *normative* for formal systems, and researchers sometimes assume that evaluative criteria must, by definition, be based on such models. This is hard to reconcile, however, with the widespread use of empirically based warrants for prescription, e.g., wherever people learn through apprenticeship with more experienced colleagues – and even more broadly, wherever it has been discovered that certain decision processes achieve objectives more reliably than others.

In this article, we will compare analytically and empirically based prescription in both theory and practice. We will argue that formal normative models are *not unique* as warrants for prescriptive intervention, and as a result the way is open for cognitive engineers to explore alternative, empirically based avenues of prescriptive research and intervention. In addition, formal normative models are not sufficiently specific to warrant actual prescriptive practice, e.g., in decision aiding, training, or consulting. Thus, cognitive engineers *must* seek alternative sources of warrant. Finally, we will describe how empirically based prescription works, respond to criticisms, and illustrate its application in systems design and the development of real-world training. We argue that empirically based normative models are a relevant and pragmatically useful source of warrant for prescriptive intervention – without denying the merits of analytically-based interventions, such as decision analysis, under suitable conditions.

Analytically Based Prescription

Von Winterfeldt (1999, p. 134) specifies four steps for analytically based cognitive engineering: (a) Identify a normative model for a judgment or decision making task; (b)

identify obstacles in the way of implementing the normative model (e.g., systematic errors or cognitive illusions in unaided decision behavior; Kahneman, Slovic, and Tversky, 1982); (c) develop tools to overcome the implementation problems; (d) test the tools and implement them if successful. These steps highlight the independence of normative modeling (step a) from description (step b) and the role of both in prescriptive modeling (step c). It also makes it clear that description has no independent justification; its purpose is simply to identify deviations from the normative model. The following discussion will focus on decision analytic consulting, but the issues we discuss are also likely to arise in training and decision aiding applications.

Decision analysis offers a repertoire of techniques for modeling choice among options with uncertain outcomes, choice among options that differ on multiple evaluative dimensions, multi-level inference from evidence to uncertain conclusions, allocation of scarce resources, and negotiation (Keeney & Raiffa, 1976; Raiffa, 1968). A decision analysis for choice under uncertainty might include the following steps: (a) Modeling the problem, which includes generating mutually exclusive options and identifying a mutually exclusive and exhaustive set of possible outcomes of the options; (b) eliciting inputs, which includes assessing the probability that each outcome will occur conditional on each of the options and the utility or degree of preference of each option-outcome combination; (c) calculating the conclusion, i.e., mathematically aggregating the probabilities and utilities for each option to calculate its subjectively expected utility; (d) testing and improving robustness, which includes an optional sensitivity analysis and iterations of modeling and elicitation steps to firm up parts of the model that are found to have a strong influence on the conclusion; and (e) communicating results, i.e., displaying and explaining the conclusion in such a way that decision makers are willing and able to select the option with the highest subjectively

expected utility (SEU). A similar process can be used to create a Bayesian inference model (Schum, 1994) or a multiattribute utility model (Keeney & Raiffa, 1976).

The role of decision theory (i.e., the SEU model) is to provide an analytical normative basis for the prescriptive methods just described. But to what extent does it succeed? We will pursue this question in two parts. First, in what sense is decision theory itself uniquely normatively justified? Second, how much does decision theory dictate (hence, warrant) actual decision analytic practice? In both respects it turns out that the analytically based warrant for decision analytic methods is less than meets the eye. Decision analysis, like empirically based prescription, must ultimately be warranted by descriptive considerations.

Is SEU Uniquely Normative?

Proponents often state that subjectively expected utility (SEU) theory uniquely defines rationality (e.g., von Winterfeldt & Edwards, 1986, p. 19). Formulation in the language of logic and mathematics sometimes lends SEU a false air of inevitability. As one commentator said, “Since mathematics [of probability and statistics] is true... we must accept the mathematics” (Rubin, 1988, p. 293). Frequent reference in the literature to “logical consistency” for compliance with SEU constraints and “logical inconsistency” for violations adds to this impression. So does misleading language that can be found in some of the normative arguments for SEU. According to Skyrms (cited by Kaplan, 1996, p. 159), for example, acceptance of SEU is necessary to avoid the “literal inconsistency” of deciding differently in the same gambling situation based merely on how the gamble is described.

These characterizations are not correct, however. When SEU is violated, what is in question is not compliance with standard systems of logic or mathematics but with specific axioms, or alleged first principles, of decision making, which have been adopted without logical or mathematical proof. In particular, the situations in which people “ought” to behave in the same way are not literally the same. They are *different* situations that are *analyzed* as

the same from the perspective of the SEU theory, which stipulates that only the probabilities and utilities of ultimate consequences “should” matter to a decision maker (Schick, 1997; Kaplan, 1996). Slovic’s (1982, p. 172) defense of SEU identifies the issue more fairly: “Maximization of expected utility commands respect as a guideline for wise behavior because it is deduced from axiomatic principles that presumably would be accepted by any rational person.”

The analytically based warrant for prescription, then, depends on the normative status of the axioms. The traditional (and perhaps still the most common) justification for them is that they appear compelling to decision makers after considered reflection, even if the same individuals do not initially behave in accordance with them. In his classic foundational book about the SEU model, Savage (1954, p. 7) argued for its normative status by asking the reader to judge if the model’s axioms were compelling:

I am about to build up a highly idealized theory of the behavior of a “rational” person with respect to decisions. In doing so I will, of course, have to ask you to agree with me that such and such maxims of behavior are “rational”.... So, when certain maxims are presented for your consideration, you must ask yourself whether you try to behave in accordance with them, or, to put it differently, how you would react if you noticed yourself violating them.

A variant of Savage’s idea does not treat intuitions as the last word on rationality, but is based instead on the idea of a *reflective equilibrium* (or coherence) between intuitions and behavior (Goodman, 1965). Individuals adjust their decision making behavior to fit principles they find intuitively compelling, but they may also adjust their principles to fit behavior they regard as rational. A normative model is justified for an individual if it reflects an equilibrium point in such a process of mutual adjustment.

MacCrimmon (1968) put Savage's proposition to a test, asking upper-middle level executives to solve problems with known decision theoretic solutions and also to indicate the extent to which they endorsed Savage's axioms. After being shown the discrepancies between the axioms and their behavior and hearing a rationale for the axioms, many of the participants wanted to conform their behavior to the axioms. However, MacCrimmon did not test whether SEU axioms were different from other principles of decision making in this regard, and thus failed to establish their unique normative status. Slovic & Tversky (1974) repeated MacCrimmon's experiment but provided rationales for either the acceptance of Savage's axioms or their rejection, and sometimes for both. The result was that decision makers were equally likely to be persuaded in the "normative" or the "non-normative" direction. Slovic and Tversky concluded that both concrete decision making behaviors *and* "uneducated" abstract judgments were irrational. In doing so, they simply assumed the normative status of the axioms, which the experiments were originally meant to test. By Savage's standard, as well as in terms of reflective equilibrium, decision makers who consistently endorse both behavior and principles that conflict with SEU axioms would have to be regarded as rational (e.g., L. J. Cohen, 1981).

In sum, MacCrimmon's and Slovic & Tversky's (1974) findings raise doubt about the unique normative status of Decision Theory. In the next section, granting that Decision Theory has normative status in some sense, we consider the extent to which it warrants, or accounts for, the success of, Decision Analysis. We will observe that Decision Theory provides very little constraint on how decision analysis (or, indeed, decision making) should be conducted. Instead, it contributes a set of rigorous and precise tools for representing beliefs and preferences, whose use must be justified by independent pragmatic arguments or by direct appeal to face plausibility. That is the core insight of Bell et al.'s (1988) emphasis on prescription as an autonomous area of research.

Does SEU Warrant Decision Analysis?

In the first stage of a decision analysis for choice under uncertainty, a set of options and outcomes must be generated for inclusion in the model. Is there a *normatively compelling* warrant for the particular options and outcomes that are selected or for the manner in which they are generated? Contrary to what is sometimes said, decision analysis does not require, even in the ideal case, the generation of all feasible or relevant options (Yates, 2001).

Decision theory constrains *subjective* beliefs and preferences; it does not aim at objective optimization. Thus, it demands consideration only of options “in the purview of the decision maker” (a phrase attributed to March by Bell et al., 1988, p. 18). As Brown, Kahr, and Peterson (1974, p. 4) correctly point out, “In its most straightforward form decision analysis serves only the function of choosing between options that have already been identified...” Thus, a decision maker who acts directly on her recognition of the appropriate action (Klein, 1993) can be coherent and rational according to SEU, unless her “purview” includes beliefs that imply the existence of a better option. Similarly, the requirement that outcomes be exhaustively specified simply means that their subjective probabilities must sum to 1. To achieve this, a model may identify only a single substantive outcome classifying all other possibilities under the catch-all “anything else”; or it may use dummy outcomes (“option a was chosen”) which have probability 1.0 conditional on each option a , to which expected utility is directly assigned. Each of these devices guarantees that whatever happens will be classifiable as one of the pre-specified possibilities.

Can it be shown that a more detailed specification of options and outcomes is always within the “purview of the decision maker”? One of the pioneers of decision theory, Ramsey (1980), thought so. He claimed that every decision maker has an opinion about everything ready for the asking (the so-called *revealed preferences* approach). On his view, beliefs and preferences about *all possible events* are revealed by the prices at which decision makers

would agree to buy and sell gambles (if forced to set a single price for both). Unfortunately, by encompassing literally everything, Ramsey's ideal model provides no guidance at all for a real analysis. A decision maker can *always* be forced to name a single price for buying and selling a gamble (thus, making no allowance for uncertainty about probabilities), no matter how unreliable her knowledge of the events upon which the bet depends. But this tells us nothing about a particular decision maker's actual interests or knowledge. If the normative suggestion is merely that more detailed modeling is always better, problems of computational complexity arise that quickly exceed the capabilities of the most powerful imaginable computer (Cherniak, 1986; Simon, 1988). Finally, the revealed preferences approach assumes that people already conform to the axioms of SEU (otherwise, they would not "have" the probabilities and utilities attributed to them). On this interpretation, Ramsey undermines the normative rationale for decision analysis by taking SEU to be descriptively accurate.

There is now a fair consensus in the decision analytic community in favor of a more psychologically plausible, *constructivist* viewpoint: that well-defined beliefs and preferences are for the most part *created* by the analysis rather than *elicited* (Bell et al., 1988, pp. 21-22; von Winterfeldt & Edwards, 1986; Fischhoff, 2000). Experiments on cognitive biases and illusions show that people provide significantly different utility and probability orderings under the influence of small, supposedly inconsequential changes in wording, response modes, or context (Kahneman, et al., 1982; Kahneman & Tversky, 2000; Gilovich, et al., 2002). Firm, precise, and orderly pre-existing beliefs and preferences are the exception rather than the rule. Thus, the "purview" of the decision maker (including relevant options, events, and outcomes) is not well defined in advance but dynamically evolves during the analysis, as the problem is identified, knowledge from appropriate domain experts is modeled, and conclusions are inferred from the modeled knowledge (Fischhoff, 2000). Most decisions about model elaboration are based on informal judgments about the reliability of the

assessments made so far and the degree of understanding or insight into the problem that has been achieved. Boundaries for model construction are determined by pragmatic considerations, not by the SEU axioms. Judgments of reliability must also be called upon to decide which assessments to revise when they are inconsistent with SEU axioms (Harman, 1986; Nozick, 1993, p. 159). Savage (1954, p. 308) was right when he said, “the theory itself does not say which way back to coherence is to be chosen.”

The elimination of biases and cognitive illusions is cited as a rationale for decision analysis. According to Raiffa (1994, p. 3), the appropriate answer to the complaint that “people often don’t behave the way normative theories say they should behave” is “to teach people how they might behave better, with more coherency, using better heuristics, with less mystical behavior, etc.... [p. 3].” From a constructivist perspective, however, the claim that behavior violates normative rules presupposes well-defined and sufficiently precise probabilities and utilities, which are not likely to exist until after an analysis (or choice experiment) has constructed them (Shaffer & Tversky, 1988). From this perspective, the analysis creates the problem (decision biases) that it is intended to solve, and perhaps should not receive so much credit (Gigerenzer, 2000). The more fundamental issue, however, remains unsettled: What is the value (or source of warrant) for the process of organizing and representing beliefs and preferences by means of decision analytic tools?

Elimination of biases may not be a sufficient warrant for decision analysis, but it is surely a necessary condition of its adequacy. Decision analytically modeled probabilities and utilities must consistently and meaningfully capture the knowledge and preferences of decision makers without regard for supposedly irrelevant variations in problem formulation and context. Research in behavioral decision making tends to raise doubts about whether such consistency is either attainable or desirable (Fischhoff, 2000). Summarizing relevant literature, Kahneman and Tversky (2000, p. 487) remark that, “Biases and cognitive illusions

are not readily eliminated by knowledge or warning.” Raiffa agrees that the debiasing challenge is serious (1994, p. 8): “Some of it may require deep therapy which should go back to early years of childhood where improper heuristics and intuitions about uncertainties develop.” If the cognitive illusions associated with probabilities and utilities run as deep as that, exploration of more natural ways to structure knowledge and preferences might be in order.

A constructivist perspective admits that there may simply not be an appropriate SEU model for some decision problems, because the assumptions about relevance imposed by the Savage axioms may not always hold. Standard SEU models exclude some variables that have a persistent and fundamental influence on human choice: For example, decision makers prefer risky situations in which they have reliable knowledge about probabilities to those in which they do not, even if the situations are identical in terms of an SEU model (Ellsberg, 1988; Einhorn & Hogarth, 1985); choices are based on probabilities only when they reflect the causal impact of actions on events rather than mere correlations (Gibbard & Harper, 1988; Pearl, 2000); and people often choose either to pre-commit their future selves or, conversely, to sacrifice the preferences of their present selves in favor of anticipated future values (Schick, 1997; McClennen, 1990). The importance people assign to knowledge, causality, and future selves seems quintessentially rational (Nozick, 1993), yet actions based on such reasons violate the SEU axioms, which require that choices depend solely on the subjective probabilities and utilities of consequences.

More generally, cognitive illusions may be accounted for by the recognition that human decision making is fundamentally *reason-based* (Lipshitz, 1994; Shafir, Simonson, & Tversky, 2000). That is, people look for salient, persuasive reasons for a decision and tend to choose the option for which such a reason can be found or constructed (Schick, 1997; Montgomery, 1993). Reasons often represent partial perspectives adopted under the influence

of a particular context and framing of the problem (Fischhoff, 2000; Schick, 1997). Apparent errors or biases will occur if probabilities and utilities reflect reasons that are active on a particular occasion rather than all the relevant information, as SEU models demand (Slovic, 2000, p. 49; Shafir, et al., 2000). One of the biggest challenges for *any* prescriptive technique, whether naturalistic decision making or decision analysis, is to find ways to capture not only the momentary reasons for decisions, but a fuller picture of the knowledge that reasons sample.

The constructivist viewpoint has elements in common with naturalistic decision making. Unlike the revealed preferences framework, constructivists recognize the importance of cognitive compatibility between models and pre-existing tendencies to organize, represent, and use knowledge. Constructivists also assume that reliable knowledge can and should be discriminated from unreliable knowledge, that solutions should be based on reliable knowledge, and that substantive (rather than logical) expertise is generally a source of reliable knowledge. On the constructivist view, there is no normative obligation to model within an SEU framework, even though it affords a formidable set of tools. *Which* tools are used and *how* they are used must be left to judgments of the analyst and client, and involve implicit empirical assumptions about the character and reliability of cognitive structures and strategies (Shafer & Tversky, 1988; Brown & Ulvila, 1976).

Other Sources of Warrant for Decision Analysis

Sophisticated decision analysts reverse the priority of pragmatics and formal justification. The success of decision analysis tends to justify the axioms underlying decision analytic tools rather than the other way around. The analytical framework serves as a means, not an end in itself. The primary sources of warrant lie in desirable features of the decision analytic process or in its likely external consequences.

Internal characteristics of the process. One warrant for decision analysis is its face validity as a *process*, i.e., its compatibility with general, common-sense intuitions regarding “good” or “rational” thinking. This includes, for example, proceeding in a deliberate, systematic fashion (“The art of good decision making lies in systematic thinking”: Hammond et al., 1998, p. 217). Improved insight and understanding of the problem are also relevant to face validity. Keeney (1999), for example, asserts that qualitative organization of thought is a more important contribution than quantitative analysis, identifying the problem is a more important function than solving it, and making good decision makers is more important than making a good decision. Many decision analysts (e.g., Hammond et al., 1998, p. 235) appear to assume, without much argument, that decision analytic strategies and techniques are *necessary* for improving understanding. Hammond, for example, points out that decision analysis helps decision makers work on the right problems, specify their objectives, create imaginative alternatives, understand the consequences, grapple with value tradeoffs, clarify their uncertainties, think hard about their risk tolerance, and consider interdependence among decisions. Face validity is also enhanced by helping decision makers avoid certain logical fallacies and decision biases that are alleged to be undesirable on their face, without need of further argument in terms of normative models (Bazerman, 1994; Kleinmuntz, Schoemaker, & Kunreuther, 1993; Russo & Schoemaker, 1987; Hammond et al., 1998). Keeney, for example, states that he now places less emphasis on formal modeling and more emphasis on avoiding decision traps that he feels are inherent in “common sense.” The paradigm of systematic, deliberative, and error-free decision making is tantamount to an informal normative model. It supplants formal consistency with SEU axioms as the motivation for modeling, although its proponents sometimes try to smuggle in features of the old paradigm, such as normative uniqueness and “logical” fallacies.

External consequences of the process. Another warrant for decision analysis is its pragmatic usefulness for achieving goals beyond the context of the analysis itself (Pitz, 1992, p. 313). Insight and improved understanding might fall under this category as well, since they extend to future problems in the same class. Objectives that refer unambiguously to external outcomes, however, include enhanced communication of a decision to others, more effective justification of the decision, and, of course, better decision outcomes. Aspects of analysis driven by external pragmatic goals are loosely connected, at best, to the original warrant of formal consistency. Ward Edwards (1983) understood this early on when he endorsed “a simple view of what the function of decision-analytic tools is: to help a decision maker with a set of problems... Decision analytic tools are *not* intended primarily for...embodying axiomatic or methodological rigor, or conforming to axiomatic structures.”

In sum, the tools for applied decision theoretic modeling are distinctively precise and robust, but their normative warrant does not share these characteristics. The most convincing rationale for decision analysis is face validity of the process and pragmatic usefulness for clients. Face validity, of course, depends on intuitions and might be characterized as a convergence (or “reflective equilibrium”) that includes the educated judgments of decision analysts themselves (influenced, of course, by decision theoretically oriented research), the subjective agreement of clients that an analysis has been illuminating or successful (influenced, of course, by persuasive efforts of the analyst), and clients’ actual acceptance of the results, as indicated by the actions they take after the analysis. Convergence among these indicators is as much an empirical question as the second rationale for decision analytic modeling: the reliable delivery of desired outcomes for clients. Neither question has been decisively answered, primarily because experimentally controlled, real-world validation data are hard to obtain, not only for decision analysis but for any other applied prescriptive technique. As Bell et al. remarked in 1988, little evidence is available to substantiate claims

of pragmatic usefulness. Von Winterfeldt (1999, p. 148) echoed a similar point more recently:

Decision analysts cannot provide much evidence of validity, and they usually emphasize the intangible benefits of decision analysis: insights and clarity, accountability and defensibility, and client satisfaction. Most decision analysts would be much more comfortable if they could point to a few excellent validation studies.

Decision analysis might eventually be validated in this way. Our point is simply that the presumption of normative uniqueness that still clings to it would nevertheless be untenable. Given its reliance on empirical warrant, its inability to represent certain decision situations, and the instability of assessments over time, why should we suppose that decision analysis is the *only* technique that might enjoy face validity and pragmatic success?

Empirically Based Prescription

Analytically based normative models presuppose a sharp distinction between errors attributable to irrationality and errors attributable to lack of knowledge or cognitive skill. From this “internalist” perspective (Plantinga, 1993, pp. 3-29), decision makers have conscious access to all the required ingredients of rational behavior – subjective probabilities and utilities, and “self-evident” intuitions about axioms. Rationality is just a matter of properly combining these pre-existing contents of their own minds. We have found that the internalist framework does not hold up: Both the justification of the axioms and the application of analytically based prescriptive tools are open empirical issues.

Empirically based prescription introduces a third-person point of view, aiming not at consistency of mental contents but at success in the real world. Two sources of empirical warrant correspond to different ways of defining or estimating such success: by reference to expert judgments and behavior, on the one hand, and direct association with desirable consequences, on the other. Associations between strategies and real-world consequences can

be either measured or inferred within the framework of ordinary scientific inquiry. A decision may be pragmatically successful, as seen or inferred from a *third person* point of view, whether or not the decision maker herself can articulate the reasons for the decision or describe the process of making it (Stich, 1991; Hammond, 1996). In some domains, conformity with norms exemplified in expert behavior or judgments may, arguably, be constitutive of success (e.g., in the arts). In domains where prescriptive intervention is of interest, however, expert behavior and judgment tend to be fallible indicators of successful outcomes.

Does Empirically Based Warrant Make Sense?

The best way to illuminate empirically based prescription may be to address some misconceptions and objections that are commonly associated with it. To start with, decision theorists accuse it of ignoring the distinction between good decisions and good outcomes. They correctly point out that outcomes are imperfect guides to decision quality under uncertainty; chance or luck may supply good outcomes for a poor decision or ruin a good decision. Empirically based prescription, however, does not apply pragmatic criteria directly to the outcomes of the individual decision to be evaluated. What counts for a sophisticated pragmatist (Rescher, 1977; Stich, 1991) is the *general* success of the *method* used to produce a belief or decision on a particular occasion. A decision has empirically based warrant to the extent that the method employed in the decision generally results in desirable outcomes under comparable conditions – regardless of the actual consequences on that occasion. Thus, formal consistency is not the only (or even the most plausible) way to distinguish good decisions from good outcomes.

A second misunderstanding is that empirically based prescription entails uncritical acceptance of human decision making as free of systematic errors (Cohen, 1993). Confusion here may arise from the fact that empirically and analytically based prescription appeal to

different benchmarks when they evaluate human decision making. The reference point for analytically based evaluation is the “ideal decision maker.” Since the ideal of perfect consistency with SEU axioms may be neither feasible nor well-defined, the analytically based researcher is virtually guaranteed to find significant irrationality in ordinary behavior and to be discouraged at the difficulty of eradicating it. Empirically based prescription, by contrast, does not designate any decision making method as inherently ideal or as inherently defective. It supports a graded evaluation, in terms of degrees of reliability of known, feasible methods under specific types of conditions, rather than all-or-none rationality. It looks for learnable, affordable, and efficient alternatives that will consistently produce better results or be effective across a broader range of circumstances, and it identifies systematic shortcomings in current practice only *relative* to such alternatives. Empirically based cognitive engineers welcome incremental improvements over the status quo – for example, novices’ approximating earlier and more closely to the performance levels of experts. Since there is no fixed ideal (such as consistency with SEU axioms) there is no obvious upper bound on the improvement of cognitive strategies. For example, in naturalistic decision making work, benchmark performance may itself continue improving as experts reflect on, share, and adaptively modify their strategies to incorporate new knowledge or meet new contingencies.

A third, related objection to empirically based prescription is that it commits the naturalistic fallacy, i.e., that it infers what is from what ought to be. But NDM researchers do not imagine that normative rules can be indiscriminately read off actual practice. On the contrary, cogent arguments supported by empirical evidence must show why (and under what conditions) the behavior in question is worthy of emulation. In fact, there are plausible arguments that link training of NDM-based expert strategies to desirable consequences: Because expert cognitive strategies represent accumulated knowledge, or “the intelligence captured by accumulated experience” (March & Heath, 1994, p. 221), it is plausible to

suppose that they are more likely to achieve success in the relevant domain than non-expert strategies. Therefore, non-experts should be more likely to achieve their own goals if they use the same (or relevantly similar) strategies as experts under comparable circumstances.

Feasibility and cost-effectiveness are already assured since the strategies derive from successful decision making practice. Since experts have, by definition, acquired them, the strategies are transferable at least in principle, although it may take novices some time to achieve the same facility.

Constructing an expert model deviates from dust-bin empiricism in another way.

Analysts do not mechanically transcribe observed behaviors, even of experts. Idealization is required to eliminate occasional performance errors, remove elements identified as irrelevant or inefficient, fill in blanks, estimate generalizability to new situations, integrate the knowledge or methods adopted by different experts, achieve theoretical coherence, and make the strategy more accessible to others. Idealization unavoidably introduces new assumptions, which complicate the prediction that a given NDM model will help non-experts achieve good consequences. Nevertheless, some empirical support exists for both the empirical success and transferability of appropriately idealized expert strategies. For example, Calderwood, Klein, & Crandall (1988) found that the first responses considered by chess masters tended to be high quality. Cohen and his colleagues (Cohen, Freeman, & Thompson, 1998; Cohen & Thompson, 2001; Cohen, 2002) found that after NDM-based training, decision strategies more nearly resembled the strategies used by experts and decisions more nearly resembled the decisions of experts.

A fourth claim is that empirical warrants for prescription are not normative, while analytically based warrants are. On the contrary, standard usage points to a more inclusive sense of *normative* than the purely formal one stipulated by Bell et al. (1988). Expression in logical or mathematical language has no bearing whatsoever on whether the guidance

provided by a model is normative (except, of course, inside the fields of logic and mathematics). *Any* set of principles, procedures, rules, etc. that is intentionally designed to give reasonable guidance about what to do is normative within the relevant context.

Normative models may utilize descriptive research, mathematical calculation, deductive proof, plausible argument, intuitive judgment, or a combination; each of these methods might be more trustworthy than the others under appropriate circumstances. As we saw earlier, reflective equilibrium arguments for decision theory combine formal deduction of rules from axioms with experimental confirmation that the axioms are intuitively acceptable to decision makers. Similarly, pragmatic justifications may combine empirical investigation of expert judgment and behavior, plausible idealization of observed strategies, and inference or experimentation to determine the reliability of the link between strategies and consequences.

A more inclusive view of normative modeling does not load the deck against decision analysis. (By contrast, the analytical paradigm keeps naturalistic decision making out of the game by *defining* rationality narrowly, in terms of explicitly formal models.) Decision analysis and naturalistic decision making can compete for empirically based warrant on a level playing field. Empirically based warrant depends on the existence of a connection between a decision making method and satisfactory outcomes, and as already noted, the case for such a connection can be made by argument or by experiment. The case for decision analysis depends on the plausibility of testable assumptions. First, it assumes that internal consistency is associated with external reliability, or more specifically, that a decision maker's chance of achieving *satisfactory* outcomes is increased by choosing the *best* action according to an explicit SEU model of beliefs and preferences. Satisfaction of this condition implies the acceptability of SEU relevance criteria (i.e., only probabilities and utilities of consequences count), the ability to provide reliable probabilities and utilities, and the absence of any other reliable method for making the decision. Second, it assumes that the deliberation

necessary to construct an SEU model is feasible and its costs are justified. As discussed earlier, there is no reason to suppose that these assumptions are always or even frequently satisfied, but when they are, deliberating about the SEU-consistency of beliefs and preferences may help an agent achieve her goals.

In sum, empirically based prescription refers to the reliability of decision making methods for producing desirable outcomes. It does not depend directly on the outcome of the decision in question, imply that decision makers are immune to systematic errors, impose a fixed benchmark of quality performance, commit the naturalistic fallacy, fail to qualify as normative, uncritically transcribe observed expert behavior, or unfairly tip the scales against decision analysis. It assures that the benchmark for improved performance is feasible, that normative judgments are relative to such a benchmark, that the benchmark dynamically changes with state-of-the-art expert performance, and that the warrant for a decision making method depends on testable empirical assumptions. None of these features are objectionable for an engineering discipline.

Empirical Warrants for Naturalistic Decision Making

In schematic form, a prescriptive intervention in NDM involves five iterative steps: (a) Conduct descriptive research on the strategies that experts use and the levels of performance associated with those strategies under various conditions. Evaluate, integrate, simplify, and supplement observed expert strategies to create a normative model capable, in principle, of achieving those levels of performance. These performance levels serve as (current) benchmarks or performance objectives. (b) Conduct descriptive research on non-expert decision making and identify pragmatically significant ways in which non-expert strategies diverge from expert strategies. (c) Derive prescriptive implications by exploring the underlying reasons for the superiority of expert over non-expert performance. Integrate those reasons into a coherent overall account of the knowledge and skill components that

distinguish expert from non-expert models. (d) Based on those implications, develop a specific prescriptive intervention (such as a decision support system, training program, or consulting method) to inculcate relevant expert skills and knowledge, and to reduce the differences between expert and non-expert performance. (e) Test the success of the prescriptive intervention in achieving more expert-like cognitive processes and improved performance. If necessary, repeat earlier steps and retest in light of the results.

Klein (1997) outlines a strategy for developing NDM-based training programs. There are two phases: A diagnostic task to identify the skills that experienced decision makers employ in a particular domain, and a design task to devise a training program that facilitates the acquisition of these skills. The diagnostic phase corresponds to steps (a), (b), and (c) in the previous paragraph and can draw on research findings regarding contrasting strategies used by experts and novices across a wide range of domains. This phase will involve a *cognitive task analysis*, or analysis of decision requirements, for the specific domain in question, to pinpoint the “difficult, critical, and/or frequent decisions that [decision makers] must make under conditions of time pressure, ambiguity, shifting situation dynamics, ill-defined goals, and other features of naturalistic environments” (Klein, et al., 1996, p. 1). Klein (1997) lists the following skills that have been identified in this manner: recognizing patterns (situation awareness), making fine perceptual discriminations, recognizing typicality and detecting anomalies, mentally simulating future states (to evaluate courses of action) and past states (to generate explanations for events) and improvising (to adapt habitual practices and standard procedures to the requirements of situations at hand). The second, design phase corresponds to steps (d) and (e) in the previous paragraph and centers on the construction and evaluation of training exercises that provide repeated cycles of practice and corrective feedback. *Deliberate practice*, with explicit goals and evaluative criteria, is a key causal factor in the acquisition of expertise (Ericsson, in press). Examples of the application of this method can

be found in the work of a number of researchers (e.g., Klein & Wolf, 1995; Nutt, 1999). An example from Cohen and his colleagues is discussed below.

A shortcoming of NDM-based training compared to SEU training is its apparent domain-specificity. If training is confined to a specific set of cues, patterns, and strategies, it is unlikely to transfer to new domains. Klein (1997) notes several related pragmatic limitations: It is sometimes difficult to identify specialized domain experts or gain access to them; the task of identifying the critical cues, patterns, and strategies can be complex; teaching them can be labor-intensive; and finally, the training will be of little use if the job changes or the person changes jobs. An alternative to directly training decision makers to think like experts is training them to *learn* like experts (Klein, 1997). In this approach the diagnosis phase focuses on the strategies that experts use to *acquire expertise*, and the design phase focuses on devising exercises for practicing those strategies. The expert learning strategies identified from the research literature by Klein (1997, p. 347) include “engaging in deliberate practice, so that each opportunity for practice has a goal and evaluation criteria; using attentional control exercises to practice flexibility in scanning situations: sampling alternative task strategies; compiling an extensive experience bank; obtaining feedback that is accurate and diagnostic and reasonably timely; enriching experiences (i.e. reviewing prior experiences to derive new insights and lessons from mistakes); building mental models; and obtaining coaching.” Another, related approach is to train the metacognitive skills by means of which decision makers monitor their own thinking and take corrective steps when problems are found (Cohen, Freeman, & Wolf, 1996; Cohen, Freeman & Thompson, 1998; Cohen & Thompson, 2001). Such skills may be more general across different domains than the specific cues, patterns, and strategies that compose specialized expertise (Means et al., 1993; Kuhn, Amsel, & O’Loughlin, 1988). Cohen and his colleagues explored computational models in which self-regulation could be shown to facilitate complex learning (Cohen, Thompson,

Adelman, Bresnick, Shastri, & Riedel, 2000b; Thompson & Cohen, 1999). Training of both expert learning skills and metacognitive skills may transfer more readily to new jobs and even new domains than job-specific knowledge.

Justification in terms of actual consequences marks a critical difference between analytically based and empirically based normative frameworks. However, the normative status of expert strategies is also supported by intuitions about face validity comparable to those invoked in support of decision analysis. Justification by consequences and justification by face validity are unlikely to be completely independent. Pragmatic successes are among the factors that shape both intuitions and actual behavior, and the degree of influence should increase as experiences accumulate (Stanovich, 1999); conversely, the pragmatic benefits of NDM training are probably enhanced by the intuitive acceptability and relative familiarity of the strategies that are trained. Indeed, trainees frequently state that they already use the strategies or close variants, despite the fact that training in such strategies improves their performance (Cohen et al., 2000a). People “vote” with their behavior, and the fact that similar strategies are widely used by proficient decisions makers across many domains supports the claim that the strategies have staying power.

Empirically Based Prescription in Action

We will illustrate empirically based prescription in detail by tracing the steps Cohen and his colleagues followed from the development of a descriptive model with normative implications to the development, application, and testing of prescriptive models for training.

(a) Descriptive and normative modeling of expert performance

Modeling of expert performance was based on an analysis of retrospective interviews with decision makers. The data set comprised 14 interviews with active-duty naval officers representing anti-air-warfare, executive officer, and tactical action officer positions in the Combat Information Center of AEGIS cruisers (Cohen, Freeman, & Wolf, 1996; Cohen,

Freeman, & Thompson., 1998; Cohen & Thompson, 2001; Kaempf, Klein, Thordsen, & Wolf, 1996), and 25 interviews with 23 Army command staff officers in operations, planning, and intelligence at division, brigade, and battalion levels (Cohen, 2002; Cohen, Freeman, & Thompson, 1997; Cohen et al., 2000a). All interviews employed a variant of the critical decision making method, a semi-structured interview technique which elicits rich and specific factual descriptions of episodes during which challenging decisions were made (Crandall, 1989; Hoffman, Crandall, & Shadbolt, 1998; Taynor, Crandall, & Wiggins, 1987). Analysis of the 14 interviews with naval officers by Klein and his colleagues resulted in a sample of 103 instances of discrete changes in situation awareness, of which 87% were coded as pattern matching and 12% as active story building (Kaempf, et al., 1996). Cohen and his colleagues focused on situations in which uncertainty or novelty blocked the application of straightforward pattern matching.

The Recognition / Metacognition (R/M) model (Cohen, Freeman and Wolf, 1996; Cohen, Freeman, & Thompson, 1998; Cohen & Thompson, 2001) accounts for the performance of highly proficient decision makers in terms of their knowledge representations and self-regulation strategies. The R/M model receives support not only from the analyses described above, but also from its theoretical coherence and consistency with: (1) previous research on expert-novice differences in problem solving and decision making, including Klein's Recognition-Primed Decision (RPD) model, (2) cognitive research on the use of narrative-like knowledge structures in reasoning (Pennington & Hastie, 1993), and (3) cognitive research on the metacognitive processes that people use to regulate memory, attention, comprehension, and problem solving (Forrest-Pressley, MacKinnon, & Waller, 1985). (Another source of support, which will not be discussed here, is supplied by neurophysiological and neural computational constraints (Cohen et al, 2000b; Thompson & Cohen, 1999.))

The RPD model (Klein, 1993; 1998) effectively distinguished between two phases of decision making: (1) A situation recognition phase in which the decision maker uses critical cues to classify the situation as an instance of a known category (e.g., a high-rise fire), and retrieves goals, expectations, and actions associated with that type of situation. (2) A critical evaluation phase in which the decision maker critiques a tentatively selected option, e.g., by mentally simulating its implementation and noting potential setbacks and obstacles. While both novices and experts follow this process (Lipshitz & Adar-Pras, 2004), experts perform it more efficiently and effectively, because they have accumulated knowledge of a larger number of typical situations (e.g., Chase & Simon, 1973) and because their self-monitoring skills are more developed (Cohen, Freeman, & Wolf, 1996).

Although people tend to represent and store information in the form of stories (e.g., Schank, 1990; Neisser, 1998), it is widely supposed that narrative modes of thinking are antithetical to reasoning or rational decision making (Bruner, 1996; Schum, 1994); stories are thought of as propaganda, not truth. Pennington and Hastie (1993; Hastie, 1993), on the other hand, presented experimental evidence supporting the role of story construction as a general strategy for understanding human action and intent. A story integrates motivations, preconditions for action, intentions, actions, and outcomes in terms of their causal and intentional relationships, using general and specific knowledge accumulated with experience. Stories enable decision makers to identify gaps where important pieces of information are missing and where, as a consequence, inferences or assumptions may be necessary. Cohen and his colleagues followed by showing how stories can be employed as tools of valid argumentation and critical reasoning (Cohen, Salas, & Thompson, 2001; Cohen, Freeman, & Wolf, 1996, Cohen, Freeman, & Thompson, 1998; Cohen & Thompson, 2001). According to the R/M model, decision makers may evaluate an uncertain hypothesis indirectly, by building the most plausible story they can around that hypothesis and then evaluating the assumptions

that were required to make the story work (e.g., to fill gaps and resolve conflicting evidence). Similarly, decision makers can evaluate a course of action by constructing and evaluating stories in which the action achieves the desired outcomes.

The R/M model identifies specific processes by means of which decision makers construct, evaluate, and modify or reject stories. These processes mesh with empirical findings in reading comprehension (Baker, 1985), metacognition (Gavelek & Raphael, 1985), the development of scientific thinking skills (Kuhn, Amsel, & O'Loughlin, 1988), and problem solving. For example, the problem solving literature suggests that experts tend to spend more time than novices revising their understanding of a complex problem and validating their conclusions (Chi, Glaser, & Rees, 1982; Patel & Groen, 1991). The reading comprehension literature describes two classes of metacognitive operations employed by skilled readers, namely, evaluation of the current state of comprehension of a text (e.g., by self-testing) and regulation of that state by remedial actions (e.g., by search of the text or by reflection). The R/M model identifies processes of critiquing and correcting required to construct and evaluate stories, i.e., situation understandings and plans that account for observed events or achieve friendly objectives by means of a relatively small number of plausible assumptions.

(b) Divergences between expert and non-expert performance

Cohen and his colleagues identified differences between expert and non-expert decision making by comparative analysis of interview transcripts, supplemented by experimental tests of specific hypotheses and by findings in the problem-solving and decision-making research literature. Among their conclusions were the following: (1) Experts are more likely than novices to build stories for handling uncertainty, while novices are more likely than experts to use "checklists" (e.g., counting the number of cues for or against a hypotheses and applying a decision rule such as majority wins). (2) Experts are more likely than novices to

step back and evaluate a story and, if necessary, build an alternative story. (3) Experts are more likely than novices to take certain types of information into account. These include big picture components that affect the intentions of the actors (e.g., motivation of a country to attack, its choice of assets to use in the attack, and its choice of targets) and some subtle details, such as the platform's localization capabilities. Differences between experts and novices in their use of such information were confirmed by a separate study. Semantic dimensions were extracted from incident descriptions by means of latent semantic analysis, and discriminant analysis was applied to the results to identify a smaller set of conceptually meaningful dimensions that distinguished more and less experienced officers (Freeman, Thompson, & Cohen, 2000). (4) Experts are more likely than novices to consider the amount of time available for decision making and to adjust the time taken accordingly. This hypothesis was experimentally confirmed in a separate study of decision making by commercial airline pilots (Cohen, Adelman, Thompson, 2000). More experienced pilots adapted to differences in a context variable (fuel remaining before a diversion decision had to be made) by adjusting the amount of time they spent thinking about a problem, taking longer when time was available and finishing sooner when it was not. Less experienced pilots were equally aware of the value of the context variable, but spent the same amount of time on decision making regardless of the variable's status.

(c) Development of prescriptive implications

The next step was to identify the elements of the normative model (R/M) corresponding to significant expert-novice differences, and to characterize the knowledge and skills associated with those elements. Such knowledge and skill elements account for the superiority of expert performance and provide the objectives of prescriptive training or decision support. In characterizing the relevant knowledge, Cohen and his colleagues emphasized causally structured relationships among events, such as story templates that

describe air or sea contacts approaching with intent to attack. In characterizing the relevant skills, they emphasized a set of specific metacognitive strategies for monitoring and regulating recognitional decision making. According to Cohen, Freeman, & Wolf (1996), these included: “going beyond pattern matching in order to create plausible stories for novel situations, noticing conflicts between observations and a conclusion, elaborating a story to explain a conflicting cue rather than simply disregarding or discounting it, sensitivity to problems in explaining away too much conflicting data, attempting to generate alternative coherent stories to account for data, and a refined ability to estimate the time available for decision making.”

(d) Development of a prescriptive training intervention

Based on these training objectives, Cohen and his colleagues developed critical thinking training for military officers both in the US Navy (Cohen & Thompson, 2001; Cohen, Freeman, & Thompson, 1998) and in the US Army (Cohen, Thompson, Adelman, Bresnick, Shastri, & Riedel, 2000a; Cohen & Freeman, 1997). The Navy version of the prescriptive model was called STEPS (Story, Test, Evaluate, Plan, and Stop). The critical thinking training consisted of four modules: (1) An overview of the STEPS cycle, which consists in story-construction, testing story-based hypotheses and expectations, evaluating the results, planning against weaknesses in the story, and stopping when time or opportunity costs outweigh the benefits of further improvement in situation understanding and plans; (2) instruction on how to construct stories, beginning with a hostile intent story template and generalizing to stories of other kinds; (3) strategies for finding and correcting problems with stories, i.e., identifying gaps in stories, looking for evidence that conflicts with story-based expectations, adopting assumptions (if necessary) to fill gaps and explain conflict, stepping back to evaluate the total package of assumptions required by the story, and modifying the

plan or adding contingencies to mitigate problems; and (4) rapidly determining whether critical thinking is appropriate and when it should stop in order to take immediate action.

(e) Testing the training

The pragmatic utility of this training has now been tested in five different studies: four studies with active-duty Navy officers and one with active-duty Army officers. All five studies demonstrated statistically significant increases in the frequency of expert-like decision making processes (such as constructing stories, identifying conflicting evidence, evaluating assumptions, and generating alternative hypotheses) as well as in the agreement of trainees' assessments and actions with those of subject matter experts. The results suggest that meta-recognitional skills can be taught effectively and applied in real tasks, and will produce decisions that are more consistent with the decisions made by experts.

Why the R/M Model Is Normative

The R/M model is a candidate for normative status with respect to time-constrained decision making in uncertain or novel situations. It represents, in an idealized form, elements that appear regularly in descriptive accounts of tactical battlefield decision making by experienced military officers. (The same type of warrant exists for Klein's RPD model, with which the R/M model is consistent.) This source of justification for R/M depends on the arguments and assumptions discussed earlier: Decision makers develop reliable decision strategies and knowledge structures over experience in a domain, and such strategies and knowledge structures can be taught to others. There is considerable evidence bearing on teachability, as discussed above. However, the link between expert performance and reliable achievement of successful outcomes has been insufficiently investigated. Experimental evidence in some domains supports the link (e.g., Calderwood, Klein, & Crandall, 1988; Chase & Simon, 1973), while other evidence does not (e.g., Shanteau, 1992). More work is

required to delineate the conditions under which expert performance is superior (Yates, 2001).

A second source of normative warrant for R/M is its face validity, or correspondence with plausible intuitions about good decision making. The face validity of decision analysis derives, we saw, from a stereotype of formal, deliberative decision making, of using logical tools to derive conclusions and evaluate consistency. The face validity of the R/M model, by contrast, rests on its use of *natural* modes of thinking, which have actually evolved in expert practice. Nevertheless, story-based strategies have been subjected to criticisms intended to refute any face validity they might possess. We argue, on the contrary, that the face validity of the R/M model is enhanced by the way it meets these criticisms.

One objection is that story-based reasoning suppresses uncertainty. Yates (2001, p. 30) points out that stories are deterministic and predicts that story-based reasoning will therefore lead to more extreme judgments (i.e., probabilities closer to 0 or 1) than are warranted by the evidence. The R/M model, however, shows how proficient decision makers handle uncertainty *within* a story based framework. Instead of directly modeling probabilistic relationships among variables (as in decision analysis), proficient decision makers represent uncertainty by means higher-level *annotations* of a basically deterministic story. Metacognitive critiquing processes monitor for specific patterns of strength of belief or preference among interdependent story elements. For example, one pattern corresponds to the absence of reasons for or against a conclusion or action; another pattern corresponds to the presence of reasons both for and against a conclusion or action (i.e., conflict); and another pattern (involving temporal fluctuation) reflects unreliable assumptions. If the stakes are high and time is available, metacognitive correcting processes may respond to such uncertainty patterns when they are found, for example, by searching for information to fill gaps in evidence, account for one or more pieces of conflicting evidence, or replace unreliable

assumptions. At key points in this iterative process, decision makers evaluate the story as a whole in terms of the uncertainty patterns that remain unresolved, including especially the reliability of remaining assumptions.

A second objection is that reason-based thinking tends to provide a partial perspective on the problem, leading to bias or instability of judgment as perspective changes (Lipshitz, 1994; Schick, 1997; Shafir, Simonson, & Tversky, 2000). The answer to this objection sheds light on two ways in which R/M strategies result in superior performance. First, reasons for a hypothesis are presented naturally in the form of stories, which weave together the hypothesis, current observations, prior knowledge, necessary assumptions, and expectations along a causally structured timeline. Therefore, decision makers can use their expectations regarding the components of specific types of stories (such as approach with intent to attack) to evaluate whether all relevant factors have been considered. Second, metacognitive strategies are specifically designed to bring knowledge to bear that would not otherwise have been used. When critiquing strategies identify specific types of uncertainty in stories, correcting strategies respond by shifting attention in active memory and by adopting or rejecting assumptions. These actions activate additional information in long-term memory that is likely to help resolve the identified uncertainty problem (Cohen et al., 2000b). The interplay of stories, metacognitive critiquing, and metacognitive correcting may help decision makers approximate the more complete use of information that decision analytic models presuppose but do not necessarily deliver.

A third pitfall sometimes associated with story building strategies is *confirmation bias* (Nisbett & Ross, 1980). In particular, stories might be used to explain away evidence against a favored hypothesis, causing decision makers to cling to beliefs long after they should have been abandoned. In highly uncertain situations, however, dropping a hypothesis at the first sign of trouble can lead to paralysis. In such situations *every* reasonable hypothesis is

problematic in the sense that some evidence can be found that it doesn't fit. For the correct hypothesis, however, an alternative explanation of apparently disconfirming evidence must be true. R/M uses this to its advantage, by taking conflicting evidence as an occasion for story building rather than mechanical rejection of a hypothesis, and by exploiting explanatory stories not as blinders but as tools of evaluation. After building a story that saves a hypothesis, decision makers step back and critically evaluate the story; the hypothesis is in trouble if the story is implausible and attempts to improve it fail. Empirical testing suggests that this strategy can work against the confirmation bias. Cohen et al. (1998) found that STEPS training significantly increased the attention officers paid to disconfirming evidence (both identifying it and seeking to explain it) along with the number of alternative hypotheses they took seriously.

In addition to pragmatic effectiveness and face validity, the R/M model draws normative warrant from a more general set of considerations, regarding necessary conditions for the acquisition of knowledge. R/M rejects foundationalism (the notion that conclusions must be derived by explicit reasoning from a fixed base of intuitively certain inputs, as in the traditional SEU paradigm) and instead emphasizes the fallible, iterative, self-corrective character of knowledge (Quine & Ullian, 1970; Popper, 1995). People use parts of what they know to challenge and correct other parts, until they arrive at an overall satisfactory account:

Empirical knowledge, like its sophisticated extension, science, is rational, not because it has a *foundation*, but because it is a self-correcting enterprise which can put any claim in jeopardy, though not all at once. (Sellars, 2000, p. 132).

Conclusion

Analytically-based prescription generally introduces decision makers to a decision process that is qualitatively different from the one they would otherwise have used. By contrast, empirically-based prescription helps decision makers realize more effective variants

of unaided decision processes, i.e., idealized versions of strategies used by experts in the same field, to which their own experience might have eventually led them. Despite this divide, there has been significant convergence between the two approaches over the years. We have noted in particular the increasing reliance of decision analytical prescription on common sense intuitions about deliberative reasoning and on pragmatic goals. This has brought a more flexible approach to decision analytic modeling and a shift in emphasis from quantitative comparison of options to enhanced understanding of the problem and communication of conclusions. At the same time, as we have noted, idealization of expert strategies plays a necessary role in naturalistic modeling of observed behavior, to enhance face plausibility and the coherence of prescriptive theories and recommendations. This aspect of empirically based prescription is too often neglected in practice. Thus, recommendations based on Cognitive Task Analysis (Gordon & Gill, 1997) might be strengthened by compatibility with normative models derived by careful idealization of expert performance.

Despite these aspects of practical convergence, analytically-based and empirically-based prescription remain basically distinct approaches to improving and aiding decision-making performance. Analytically-based prescription is a top-down approach, deriving the warrant for prescriptive tools from formally represented first principles. Empirically-based prescription is primarily a bottom-up approach, which derives normative models by idealization of actual behavior identified as proficient, and which derives prescriptive tools by specifying the skills and knowledge that enable experts to execute the normative model well. The R/M model is normative because it incorporates a descriptively accurate, albeit idealized, model of how experienced decision makers make time-stressed decisions under uncertainty. The skills of story construction and critical thinking that it embodies are directly translatable into a prescriptive model such as STEP.

The future development of both empirically based and analytically based prescription hinges on their respective responses to the same challenge: to demonstrate and improve their pragmatic effectiveness. Although some degree of competition is inevitable, there are also significant opportunities for each approach to support the other in achieving these goals. This will be accomplished by (a) rigorous evaluation studies that focus on success in achieving real-world objectives; (b) delineation of respective boundaries of application, i.e., decision types and contexts for which each approach is likely to be effective (Yates, 2001), without excluding the possibility of a substantial overlap; (c) development and testing of methods which exploit opportunities for integrating decision analytic and naturalistic methods where they both apply, and (d), though we have not discussed it in this article, closer attention to the similarities and differences in the dynamic and contextual aspects of advice giving and taking. With respect to the last point, decision analysis and NDM-based prescription might ultimately be best understood as species of dialogue, with characteristic roles, phases, end-states, and constraints on permissible contributions by occupants of each role at each phase (Cohen, 2004). Within the family of advice-giving exchanges, decision analysis and NDM-based prescription are alike in drawing prescriptive warrant from normative models and in providing a role for specialized consultants to construct and communicate the models. They differ in the nature of the models they use and therefore in the means consultants employ to gain clients' confidence, identify their needs, and generate solutions. The purpose of this article has been to take a step toward the goals outlined above, by increasing awareness of the normative underpinnings of the different models, by drawing attention to the warrants they provide for prescriptive intervention, and by encouraging critical reflection and debate on their nature and merits. The ultimate objective is an approach to cognitive engineering that is both more integrated and more effective.

References

- Baker, L. (1985). How do we know when we don't understand? Standards for evaluating text comprehension. In D.L. Forrest-Pressley & T.G. Waller (Eds.), *Metacognition, cognition, and human performance*, 1, (pp. 155-205). Orlando, FL: Academic Press.
- Baron, J., & Brown, R. V. (1991). *Teaching decision making to adolescents*. Mahwah, NJ: Lawrence Erlbaum.
- Bazerman, M.H. (1994). *Judgment in management decision making* (2nd ed.). New York: Wiley.
- Bell, D.E., Raiffa, H., & Tversky, A. (1988). Descriptive, normative, and prescriptive interactions in decision making. In D. E. Bell, H. Raiffa, & A. Tversky (Eds.), *Decision making: Descriptive, normative, and prescriptive interactions* (pp. 9-30). Cambridge, MA: Cambridge University.
- Brown, R.V., Kahr, A.S. & Peterson, C. (1974). *Decision analysis for the manager*. New York, NY: Holt, Rinehart and Winston.
- Brown, R. V. & Ulvila, J. W. (1976). *Selecting analysis approaches for decision situations: A matching of taxonomies*. McLean, VA: Decisions and Designs, Inc.
- Bruner, J. (1996). Frames for thinking. In D. R. Olson & N. Torrance (Eds.), *Modes of thought: Explorations in culture and cognition* (pp. 93-105). Cambridge: Cambridge University.
- Calderwood, R., Klein, G.A., & Crandall, B.W. (1988). Time pressure, skill, and move quality in chess. *American Journal of Psychology*, 101, 481-491.
- Chase, W., & Simon, H. (1973). The mind's eye in chess. In W. Chase (Ed.), *Visual Information Processing* (pp. 215-281). New York: Academic Press.

Cherniak, C. (1986). *Minimal rationality*. Cambridge, MA: MIT.

Chi, M., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R.S. Sternberg (Ed.), *Advances in the psychology of human intelligence*, 1 (pp. 7-75). Hillsdale, NJ: Erlbaum.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *The Brain and Behavioral Sciences*, 4(3), 510-517.

Cohen, M. S. The naturalistic basis of decision biases. In Klein, G.A., Orasanu, J., Calderwood, R., & Zsombok, C.E. (Eds.), *Decision making in action: Models and methods*. Norwood, NJ: Ablex Publishing Corporation, 1993.

Cohen, M. S. (2002). Training critical thinking to enhance battlefield initiative. In *Proceedings of the Human Factors and Ergonomics Society's 46th Annual Meeting*. Human Factors & Ergonomics Society.

Cohen, M. S. (2004). Leadership as the orchestration and improvisation of dialogue: Cognitive and communicative skills in conversations among leaders and subordinates. In D. V. Day, S. J. Zaccaro et al. (Eds.) *Leader development for transforming organizations: Growing leaders for tomorrow* (pp. 177-208). Mahwah, NJ: Erlbaum.

Cohen, M.S., Freeman, J.T., & Wolf, S. (1996). Meta-recognition in time stressed decision-making: Recognizing, critiquing, and correcting. *Human Factors*, 38, 206-219.

Cohen, M. S., & Freeman, J. T. (1997). Understanding and enhancing critical thinking in recognition-based decision-making. In R. Flin, E. Salas, M. Strub, & L.

Martin (Eds.), *Decision making under stress: Emerging themes and applications* (pp. 161-169). Aldershot, UK: Ashgate.

Cohen, M., Freeman, J.T., & Thompson, B. (1997). Training the naturalistic decision maker. In G. K. C. Zsombok (Ed.), *Naturalistic Decision Making*, (257-268). Mahwah, NJ: Erlbaum.

Cohen, M. S., Freeman, J. T., & Thompson, B. (1998). Critical thinking skills in tactical decision making: A model and a training strategy. In J. A. Cannon-Bowers & E. Salas (Eds.), *Decision making under stress: Implications for training and simulation* (155-189). American Psychological Association.

Cohen, M. S., Adelman, L., & Thompson, B.T. (2000). Experimental investigation of uncertainty, stakes, and time in pilot decision making. Arlington, VA: Cognitive Technologies.

Cohen, M. S., Thompson, B.B., Adelman, L., Bresnick, T.A., Shastri, L., & Riedel, S. (2000a). *Training critical thinking for the battlefield. Volume II: Training system and evaluation*. Arlington, VA: Cognitive Technologies, Inc.

Cohen, M. S., Thompson, B.B., Adelman, L., Bresnick, T.A., Shastri, L., & Riedel, S. (2000b). *Training Critical Thinking for the Battlefield. Volume III: Modeling and Simulation of Battlefield Critical Thinking*. Arlington, VA: Cognitive Technologies, Inc.

Cohen, M. S., & Thompson, B. B. (2001). Training teams to take initiative: Critical thinking in novel situations. In E. Salas (Ed.), *Advances in Cognitive Engineering and Human Performance Research, Vol. 1*, (pp. 251-291). Amsterdam: JAI.

Cohen, M. S., Salas, E., & Thompson B.B. (2001). *Critical thinking: A normative, cognitive, and pragmatic framework*. Arlington, VA: Cognitive Technologies, Inc.

Crandall, B., & Getchell-Reiter, K. (1993). Critical decision method: A technique for eliciting concrete assessment indicators from the intuition of NICU nurses. *Advances in Nursing Science*, 16 (1), 42-51.

Edwards, W. (1983). *Simplicity in decision analysis: Two case studies*. Cambridge: Harvard College.

Einhorn, H. J., & Hogarth, R. M. (1985). Ambiguity and uncertainty in probabilistic inference. *Psychological Review*, 92(4), 433-461.

Ellsberg, D. (1988). Risk, ambiguity, and the Savage axioms. In P. Gardenfors & N.-E. Sahlin (Eds.) *Decision, probability, and utility: Selected reasoning* (pp. 245-269). Cambridge: Cambridge University.

Ericsson, K.A. (2005). Superior decision making as an integral quality of expert performance: Insights into the mediating mechanisms and their acquisition through deliberate practice. In B. Brehmer, R. Lipshitz, & H. Montgomery (Eds.), *How professionals make decisions*, (135-167). Mahwah, NJ: Lawrence Erlbaum Associates.

Ericsson, K.A., & Smith, J. (1991) (Eds.), *Toward a general theory of expertise: Prospects and limits* (1-38). Cambridge: Cambridge University.

Fischhoff, B. (1991). Value elicitation: Is there anything in there? *American Psychologist*, 46, 835-847.

Forrest-Pressley, D., MacKinnon, G.E., & Waller, T.G. (1985). *Metacognition, cognition, and human performance*. New York: Academic Press.

Freeman, J. T., Thompson, B. T., & Cohen, M. S. (2000). Modeling and diagnosing domain knowledge using latent semantic indexing. *Interactive Learning Environments: Special Issue on Latent Semantic Indexing*.

Gavelek, J., & Raphael, T.E. (1985). Metacognition, instruction, and the role of questioning activities. In D.L. Forrest-Pressley, G.E. MacKinnon, & T.G. Waller (Eds.), *Metacognition, cognition, and human performance*, 2 (103-136). New York: Academic Press.

Gibbard, A. & Harper, W. L. (1988). Counterfactuals and two kinds of expected utility. In P. Gardenfors & N.-E. Sahlin (Eds.) *Decision, probability, and utility: Selected reasoning* (pp. 341-376). Cambridge: Cambridge University.

Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. Oxford: Oxford University.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge, UK: Cambridge University Press.

Goodman, N. (1965). *Fact, fiction, and forecast*. Indianapolis, IN: Bobbs-Merrill.

Gordon, S.E. and Gill, R.T. *Cognitive Task Analysis*. In Zsombok, C.E. and Klein, G. (Eds.), *Naturalistic Decision Making* (Pp. 131-140), Mahwah, NJ: Erlbaum, 1997.

Hammond, J.S., Keeney, R.L., & Raiffa, H. (1998). *Smart choices: A practical guide to making better decisions*. Cambridge, MA: Harvard University.

Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT.

Johnson, R.H., & Blair, J.A. (1994). *New essays in informal logic*. Windsor, Ontario: Vale.

Hammond, K.R. (1996). *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice*. Oxford: Oxford University.

Hastie, R. (1993). *Inside the juror*. Cambridge UK: Cambridge University Press.

Hoffman, R.R., Crandell, B., & Shadbolt, N. (1998). Use of critical decision method to elicit expert Knowledge: A case study in the methodology of expert task analysis. *Human Factors*, 40, 254-276.

Kaempf, G.F., Klein, G., Thordsen, M.L., & Wolf, S. (1996). Decision making in complex command-and-control environments. *Human Factors*, 38, 206-219.

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

Kahneman, D., & Tversky, A. (Eds.). (2000). *Choice, values, and frames*. Cambridge: Cambridge University.

Kaplan, M. (1996). *Decision theory as philosophy*. Cambridge UK: Cambridge University.

Keeney, R.L. (1999). From 'decisions with multiple objectives' to 'smart choices'. Paper presented at the 17th conference on Subjective Probability, Utility, and Decision Making, Mannheim, Germany, August 9-11.

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

Klein, G. A. (1993). Recognition-primed decision (RPD) model of rapid decision making. In G.A. Klein, J. Orasanu, R. Calderwood, & C. Zsombok (Eds.), *Decision making in action: Models and methods* (pp. 138-147). Norwood, CT: Ablex.

Klein, G. (1997). Developing expertise in decision making. *Thinking and Reasoning*, 34, 337-352.

Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT.

Klein, G., Kaempf, G.L., Thordsen, M., & Miller, T. (1996). The uses of decision requirement. *International Journal of Human Computer Studies*, 46, 1-15.

Klein, G.A., Orasanu, J., Calderwood, R., & Zsombok C. (1993). (Eds.), *Decision making in action: Models and methods*. Norwood, NJ: Ablex.

Klein, G. & Wolf, S., (1995). *Decision-centered training*. Proceedings of the 39th Annual Meeting of the Human Factors and Ergonomics Society, San Diego, CA. Human Factors and Ergonomics Society, Inc: Santa Monica, CA.

Kleinmuntz, D., Schoemaker, P.J.H., & Kunreuther, H.C. (1993). *Decision sciences*. Cambridge, MA: Cambridge University.

Kuhn, D., Amsel, E., & O'Loughlin, M. (1988). *The development of scientific thinking skills*. San Diego: Academic Press.

LeBoeuf, R., & Shafir, E. (2001). Problems and methods in naturalistic decision making research. *Journal of Behavioral Decision Making*, 14, 374-375.

Lipshitz, R., (1994). Decision making in three modes. *Journal for the Theory of Social Behavior*, 24, 47-66.

Lipshitz, R., & Adar-Pras, A. (2004). Recognition-primed decisions in the laboratory. In H. Montgomery, R. Lipshitz, & B. Brehmer, (Eds.). *How professionals make decisions* (91-106). Hillsdale, NJ: Erlbaum.

Lipshitz, R., Klein, G., Orasanu, J., & Salas, E. (2001) Focus article: Taking stock of naturalistic decision making. *Journal of Behavioral Decision Making*, 14, 331-352.

MacCrimmon, K. R. (1968). Descriptive and normative implications of the decision theory postulates. In K. Borch & J. Mossin (Eds.), *Risk and uncertainty* (pp. 3-23). New York: St Martin's.

March, J.G., & Heath, C. (1994). *A primer on decision making*. New York: Free Press.

McClennen, E.F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge: Cambridge University.

Means, B., Salas, E., Crandall, B., & Jacobs, T.O. (1993). Training decision makers for the real world. In J. O. G.A. Klein, R. Calderwood, & C.E. Zsombok (Ed.), *Decision Making in Action: Models and Methods* (pp. 306-326). Norwood, NJ: Ablex.

Montgomery, H. (1993). The search for a dominance structure in decision making: examining the evidence. In J. O. G. A. Klein, R. Calderwood, & C. E. Zsombok (Ed.), *Decision Making in Action: Models and Methods* (pp. 182-187). Norwood, NJ: Ablex.

Neisser, U. (1998). Stories, selves, and schemata: A review of ecological findings. In S. Gathercole & C. Cornoldi (Eds.), *Theories of Memory, Vol. 2*. (171-186) Hove, UK: Psychology Press.

Nisbett, R. E. (Ed.). (1993). *Rules for reasoning*. Hillsdale, NJ: Erlbaum.

Nisbett, R., & Ross, L. (1980). *Human inference: strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall, Inc.

Nozick, R. (1993). *The nature of rationality*. Princeton, NJ: Princeton University.

Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Executive*, 13, 75-90.

Patel, V.L., & Groen, G.J. (1991). The general and specific nature of medical expertise: A critical look. In K.A. Ericsson & J. Smith (Eds.), *Towards a general theory of expertise* (pp. 93-125). Cambridge, MA: Cambridge University.

Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge UK: Cambridge University.

Pennington, N., & Hastie, R. (1993). A theory of explanation-based decision making. In G.A. Klein, J. Orasanu, R. Calderwood, & C. Zsombok (Eds.), *Decision Making in Action: Models and Methods* (pp. 188-201). Norwood, CT: Ablex.

Pitz, G.F. (1992). Risk taking, design, and training. In J.F. Yates (Ed.), *Risk-taking behavior* (pp. 283-329). New York: Wiley.

Plantinga, A. (1993). *Warrant: the current debate*. 1993: Oxford University.

Pliske, R., McCloskey, R., & Klein, G. (2001). Facilitating learning from experience: An innovative approach to decision skills training. In E. Salas and G. Klein (Eds.), *Linking expertise and naturalistic decision making* (pp. 37-53). Mahwah, NJ: Erlbaum.

Popper, K. (1995). *Conjectures and refutations* (1963). New York: HarperCollins.

Quine, W. V. O., & Ullian, J. S. (1970). *The web of belief*. New York: Random House.

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Reading, MA: Addison Wesley.

Raiffa, H. (1994). The prescriptive orientation of decision making: A synthesis of decision analysis, behavioral decision making, and game theory. In S. Rios (Ed.), *Decision theory and decision analysis: Trends and challenges*, (pp. 3-13). Dordrecht: Kluwer.

Ramsey, F. (1980). Truth and probability (1931). In H. Kyburg, Jr. & H. E. Smokler (Eds.), *Studies in subjective probability*. Huntington, NY: Krieger.

Rescher, N. (1977). *Methodological pragmatism: A systems -theoretic approach to the theory of knowledge*. New York, NY: New York University.

Rubin, H. (1988). Discussion of A. P. Dempster, Probability, evidence, and judgment. In D. E. Bell, H. Raiffa et al. (Eds.) *Decision Making: Descriptive, normative, and prescriptive interactions* (pp. 292-293). Cambridge: Cambridge University.

Russo, E.J., & Schoemaker, P.J.H. (1987). *Decision traps: Ten barriers to brilliant decision making and how to overcome them*. New York: Doubleday.

Savage, L.J. (1954). *The foundations of statistics*. New York: Wiley.

Schank, R. C. (1990). *Tell me a story: Narrative and intelligence*. Evanston IL: Northwestern University.

Schick, F. (1997). *Making choices: A recasting of decision theory*. Cambridge UK: Cambridge University.

Schum, D.A. (1994). *The evidential foundations of probabilistic reasoning*. New York, NY: John Wiley & Sons.

Sellars, W. (2000/1975). *Epistemic Principles*. In E. Sosa & J. Kim (Eds.) *Epistemology: An anthology* (pp. 125-133). Malden, MA: Blackwell.

Shafer, G. & Tversky, A. (1988). Languages and designs for probability judgment. In D. E. Bell, H. Raiffa et al. (Eds.) *Decision Making: Descriptive, normative, and prescriptive interactions* (pp. 237-265). Cambridge: Cambridge University.

Shafir, E., Simonson, I., & Tversky, A. (2000). Reason-based choice. In D. Kahneman & A. Tversky (Eds.) *Choice, values, and frames* Cambridge: Cambridge University.

Shanteau, J. (1992). The psychology of experts: an alternative view. In G. Wright & F. Bolger (Eds.) *Expertise and Decision Support*. (11-23). New York: Plenum.

Simon, H. (1988). Rationality as process and as product of thought. In D. E. Bell, H. Raiffa et al. (Eds.) *Decision Making: Descriptive, normative, and prescriptive interactions* (pp. 58-77). Cambridge: Cambridge University.

Skyrms, B. (1984). *Pragmatics and empiricism*. New Haven, CT: Yale University Press.

Slovic, P. (1982). Toward understanding and improving decisions. In M.D. Dunnette & E.A. Fleishman (Eds.), *Human capability assessment*. Hillsdale, NJ: Erlbaum.

Slovic, P. (2000). The construction of preference. In D. Kahneman & A. Tversky (Eds.) *Choice, values, and frames*. (489-502). Cambridge: Cambridge University.

Slovic, P., & Tversky, A. (1974). Who accepts Savage's axioms? *Behavioral Science*, 19(6), 368-373.

Stich, S. (1991). *The fragmentation of reason*. Cambridge, MA: MIT.

Thompson, B. B., & Cohen, M. S. (1999). Naturalistic decision making and models of computational intelligence. In A. Jagota, T. Plate, L. Shastri, & R. Sun (Eds.), *Connectionist symbol processing: Dead or alive? Neural Computing Surveys*, 2, (1) 26-28.

Taynor, J., Crandall, B., & Wiggins, S. (1987). *The reliability of the critical decision method*. Contract MDA-903-C-0170. Alexandria, VA: U.S. Army Research Institute Field Unit.

Thompson, B. B., & Cohen, M. S. (1999). Naturalistic decision making and models of computational intelligence. In A. Jagota, T. Plate, L. Shastri, & R. Sun (Collective article Eds.), *Connectionist symbol processing: Dead or alive?* *Neural Computing Surveys*, 2, 26–28.

von Winterfeldt, D. (1999). On the relevance of behavioral decision research for decision analysis. In J. Shanteau, B. A. Mellers et al. (Eds.) *Decision science and technology: Reflections on the contributions of Ward Edwards*, (133 – 154). Dordrecht, The Netherlands: Kluwer Academic Publishers.

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

Yates, J. F. (2001). "Outside" impressions of naturalistic decision making. In E. Salas & G. Klein (Eds.), *Linking expertise and naturalistic decision making* (pp. 9-33). Mahwah, NJ: Erlbaum.

Biographical notes

Raanan Lipshitz is an Associate Professor at the Department of Psychology, the University of Haifa, Haifa, Israel. He obtained his Ph.D. degree in Administrative Behavior from Purdue University in 1976.

Marvin Cohen is President and Chief Scientist of Cognitive Technologies, Arlington, Virginia. He obtained his Ph.D. degree in Experimental Psychology from Harvard University in 1980.