FOUNDATIONS FOR A MODEL OF KNOWING: I. CONSTRUCTING REALITY

Neil M. Agnew
York University

John L. Brown
University of Alberta

ABSTRACT

Traditional views of knowledge are being challenged. An emerging "constructivist" perspective, as proposed by George Kelly, an engineer turned clinician, suggests that to a large degree we construct reality. In his "constructive alternativism" Kelly assumes that we validate our hypotheses and beliefs through *subjectively construed* goodness-of-fit criteria applied to perceived differences between anticipations and feedback. His model of construing is compatible with those emerging in the history and philosophy of science and in cognitive psychology. Nevertheless, constructivists must answer a perplexing question: How can fallible knowledge, constructed as it is from abstracted and incomplete representations of objects and events, capture and maintain our confidence, as it does, and furthermore prove highly functional, as it does?

> Much or even most of the work in philosophy of science today which is not closely tied to specific historical or current scientific case studies embodies a metaphysical stance which, in effect, assumes that the scientist is an omniscient and computationally omnipotent Laplacean demon. (Wimsatt, 1981, pp. 151-152)

Most social scientists take "knowledge" for granted, believe that science, properly conducted, produces the highest form of knowledge, the "fact," and that scientific theories summarize and predict facts.

Even when acknowledging the relevance of Plato's metaphor of the cave, or of paradigm revolutions, or of experimenter effects, most of us view such historical oddities as exceptions, or as arising from naive misperceptions afflicting others, or as correctable by serious investigators. Moreover, we see ourselves as relatively immune from such distortions, and so somehow able to penetrate the iron curtain of culture and confront reality nose-to-nose. Except during moments of personal angst, subjectivity, like determinism, is a condition of "the other." For example, I see the behaviour of my experimental subjects as determined and decodable. Others are afflicted with cultural, or theoretical, or methodological blinders, not me. When momentarily we do manage to "step ouside" our particular frame of reference, like superior anthropologists, we can briefly entertain conflicting constructions of reality.

How can we reconcile the compelling yet incompatible "realities" of hard-nosed empiricism and social constructivism? One popular and remarkably functional solution involves believing a given ontological/epistemological framework — realism or constructivism — and proceeding with the accumulation of positive evidence for one's particular hypothesis or theory, while avoiding or "managing" negative or anomalous data. While, from the perspective of a bounded rationality, this is probably what we do most of the time, it is not always a satisfactory solution. As scientists, whether realists or constructivists, we may periodically become aware that Boring's Zeitgeists might actually afflict us as well as historical others, or of the data diddling of a Sir Cyril Burt, or that a significant portion of the

We acknowledge the assistance of Fred Weizmann of York University, a former student of Kelly's who shared reference material that would otherwise have been unavailable. Some of the ideas expressed in this paper were presented at the Second Biennial Conference of the North American Personal Construct Network, at Banff in June, 1986, and at the American Psychological Association meetings in Washington in August, 1986. We have greatly benefited from the writings of, and discussions with, Mildred Shaw, Brian Gaines, Jack Adams-Webber, Michael Mahoney and Ken Ford. Finally we wish to sincerely acknowledge the very helpful comments of the anonymous reviewers of this article.

research literature consists of culled type 1 errors, or of the strong suspicion that a hidden or excluded interaction term could negate a previously significant ANOVA, or of the awesome possibility that Kuhn's notion of "scientific revolutions" is not merely discounting some archaic beliefs, but rather threatening the very integrity of SCIENCE. Happily such distressing flashes are transitory, leaving us free to go about our scientific business without getting lost in the infinite regress of ontological or epistemological assumptions on which our empirical, or constructivist theories rest.

Many philosophers and a few psychologists have risked cognitive burnout by exploring the presumptive underpinnings of our "knowledge." In this article we explore the question: "What do you know for sure?" concerning a reality which seems for the most part exceedingly complex and intricately interactive — at least that is our ontological assumption. In developing a model for knowing such an unobservable and labyrinthine reality, we will rely on various authors, but particularly on the work of George Kelly (1955), a clinician, and Donald Campbell (1977, 1987a, 1987b), a social psychologist and highly respected methodologist. They take seriously the arguments both of the "realists," who propose we have more or less direct access to "reality," and of the relativists, who propose that we invent or construct our realities out of abstractions or "shadows." We believe that Kelly and Campbell provide the foundations for a theoretical integration of what are apparently incompatible ontological and epistemological perspectives.

We will refer to the work of other scholars as well in relation to a central Kelly/Campbell theme, namely that knowledge — personal, social and scientific — is the product of selected abstractions guided by given implicit and explicit anticipations. At first glance this theme may sound like the conceptual anchor for a constructivist manifesto; nevertheless, we will argue that it can accommodate a sophisticated and critical empiricism as well.

In this paper we focus on Kelly's constructivism, and also on supportive work by current investigators modelling ways in which we can generate and maintain confidence in constructed, and highly fallible, knowledge. In a companion paper (Agnew & Brown, 1989) we discuss the conditions under which much of our knowledge may be fallible but functional, and also explore the case for realism, focusing particularly on Campbell's (1987a, 1987b) "hypothetical realism."

Kelly's Perspective: Constructive Alternativism

Kelly's theory is a mixture of general postulates and clinical observations. Nevertheless, it lays out the conceptual foundations for developments in the philosophy, sociology, and psychology of science. As noted by Mahoney (1988a), in a scholarly review of constructive metatheory, Kelly's model of construing remains undeservedly obscure. The theory is clinically derived and has a deductive aura, in the sense that Kelly formulates postulates and corollaries. Perhaps Kelly triggers unpleasant reminders of Hull's "hypothetico-deductive" approach. In any case his model is unfamiliar to, or discounted by, investigators more accustomed to employing experimental methods and inductively accumulating positive instances. Nevertheless, we will argue that Kelly's core assumptions are highly relevant to social science in general, and increasingly reflected in modern decision theory and cognitive science.

Underlying Kelly's personal construct theory (1955) is the principle of "Constructive Alternativism." "This principle asserts that reality does not directly reveal itself to us, but rather it is subject to as many alternative ways of construing it as we ourselves can invent. Hence the variety of human experience" (Adams-Webber 1979, p 1).

This constructivist principle serves to recall the parable of Plato's Cave, where the inhabitants construct their knowledge of reality from the shadows reflected on the wall, having no direct access to objects or events causing the shadows. While, in our view, Kelly does not rule out the possibility of veridical reflections of reality (realism), he proposes that our goodness-of-fit criteria (our judgements of how well the shadows represent reality) are subjectively construed, and, by implication, that we are dealing with an open system subject to "continual revision."

More particularly, Kelly's theory provides theoretical guidelines for helping address a recurrent question arising in the history and philosophy of science and in artificial intelligence studies: How do we reduce the problem under investigation to manageable cognitive size? Herbert Simon (1977) summarizes the issue:

> Because the central nervous system can only do a few things at a time, and because the human memory and the human environment jointly contain an enormous amount of information potentially relevant to behavior, it is essential that there exist processes to determine what tiny fraction of this totality will be evoked at any given moment, and will during that moment constitute the effective environment of thought and behavior. (p. 159)

Here we see the need for tacit processes of selective attention that *automatically* and profoundly reduce any "search space," and, within that reduced space, implicitly select for isolation and/or manipulation a tiny subset — a salient or familiar family — of the potentially infinite number of interacting variables. Fodor (1983) discusses the issue of inherent limits of knowing in terms of what he calls "Epistemic Boundedness: the idea that there are endogenously determined constraints on the kinds of problems that human beings can solve, hence on the kinds of things we can know" (p. 120), or "that our cognitive organization [modular or non-modular] imposes epistemically significant constraints on the beliefs that we can entertain" (p. 120).

Kelly (1955) early recognized that individuals must possess mechanisms that automatically restrict their range of attention. He postulated that a construct, or a hierarchy or network of constructs, bounds our anticipations of particular experience, and selects abstractions from possible worlds (large or small), to serve the anticipations. Our contructs reflect our bounded rationality by limiting the number of events addressed, and by operating within a restricted or manageable frame of reference. Kelly, with his Range Corollary, anticipates Simon's bounded rationality theme (Simon 1977, 1979), noting that "A construct is convenient for the anticipation of a finite range of events only" (Kelly 1955, p. 68). His notion of constructs having "a limited range of convenience" also anticipates Von Wright's (1966) "range of relevance" concept, as well as the so-called "frame problem" (Adams-Webber, 1988; Ford 1988).

> Generalizations have an associated "range of relevance" and consequently, only things within a generalization's range of relevance may constitute confirming or disconfirming evidence, while all other things (things outside the range of relevance) are irrelevant. Furthermore, when a generalization's range of relevance is not specified (as is nearly always the case) it is taken to be the *natural* range of relevance, which is defined as the class of things that fall under the antecedent term. (Ford, 1989, p. 6)

Notice the cognitive burden carried by the word "natural" in the above quotation. The word implies that through genetics and/or culture we frame or close what would otherwise be open incomprehensible systems. But on what mechanisms or structures do we rely for such tacit closures? Do we rely on the *structure of reality* inherent in our genetically determined veridical feedback mechanisms — a seeing-is-believing model? Or on the *structure of beliefs* operating through socially constructed feedforward mechanisms, like anticipations, hypotheses, biases — a believing-is-seeing model? Or on some complex interaction of the two? Somehow we are able to focus our bounded rationality, our limited attention and analytic capacity on what is "relevant." But how do we know what is relevant? We face the relevance puzzle in a world in which, the physicists tell us, everything is related and so relevant. In the domain of nuclear energy when you shake one thing everything else rattles. In the domain of the upper atmosphere when you poke a hole in the ozone layer everything below starts to shrivel. But we also live in a world, physical and social scientists tell us, in which some things are more closely related and so more relevant than others. So in some instances when you shake one thing only a few other things rattle, at least rattle immediately and obviously.

Therefore, the relevance issue presents a profound puzzle. Not so much a puzzle for citizens who intuitively know most of the time what to attend to and what to ignore. To be sure there are intervals of anxiety and indecision when we can't decide what is and what is not relevant, and are puzzled at how others can be so certain, how they can really select the evidence to fit the needs of their own beliefs. But the puzzle is truly profound for cognitive and computer scientists attempting to figure out how we intuitively draw our large and small boundaries, and how effortlessly our attention selects and links relevant representations from the flow of external and internal stimuli (Pylyshyn, 1987). We don't know what automatic mechanisms we use to decide what is irrelevant and what is relevant in the flood of stimuli in different domains, ranging from passing the butter to selecting a mate.

Kelly helps address this neglected puzzle by suggesting that we rely on a limited set of

abstractions selected from the flow of experience, and that such abstractions service — are selected to be relevant to — our anticipations. Our anticipations, valid or otherwise, play a large role in determining what is considered relevant. The degree of compatibility between anticipations and abstracted feedback is determined by subjective criteria.

The puzzle of relevance underlies the "frame problem" (Fodor, 1983; pp. 112–113): "The problem of putting a 'frame' around a set of beliefs that may need to be revised in the light of specified newly available information." For example, the rejection of an article you submitted for publication may affect, or be construed as relevant to, only a tiny sub-set of your beliefs (e.g., relating to the questionable intelligence of the journal editor and his/her selection of reviewers) and trigger the decisive and immediate action of resubmitting it to another journal. Or such a rejection may resonate throughout your system of constructs by threatening a core belief (e.g., be construed as relevant to a core belief relating to your own intellectual and academic competence and linked to most of your other beliefs) and you may, therefore, be unable to stop thinking about it. In Kelly's terms, how we frame our experience is a function of subjectively construed goodness-of-fit criteria applied to the difference between our anticipations and our abstracted representations of events. Various scholars attempt to model the boundaries and internal structure or integrity of our beliefs or cognitions (Popper, 1974a & 1974b; Jones, 1976; Campbell, 1977; Fodor, 1983; Gaines & Shaw, 1984; Pylyshyn, 1985; Rychlak, 1987; Mahoney, 1988a, 1988b)[1]. We will discuss later in the paper the problems of bounding or framing experience, and of extracting functional knowledge from such radically reduced domains.

At this point we suggest that various degrees of isomorphism (ranging from high to low and/or "pathological") are possible between subjective and "external" criteria, whether the external criteria are "objective," that is scientific or Popperian in the realist sense, or cultural or tribal in the Kuhnian sense. Therefore, in our view, an acceptance of Kelly's Constructive Alternativism supplemented by the Range of Convenience corollary does not imply an exclusive commitment to ontological relativism or constructivism. It does, however, commit one to a much more sophisticated epistemology than espoused by traditional realism or positivism and its various revisions. Because of his constructivist epistemological assumptions, a cursory reading of Kelly could lead one to believe that he was not merely an epistemological constructivist who assumes our methods of mapping reality significantly influence or distort the maps we obtain, but also an ontological constructivist who assumes we go ahead willy-nilly and invent possible words (realities), any one of which is as likely to be as valid as the other because reality, like beauty, lies in the eye of the beholder. However, we interpret Kelly's theory as ontologically accommodating not only for constructivism, but also for a critical or hypothetical realism.

While not philosophically naive, Kelly (1955) devoted little time to formal ontological arguments. He states his position as follows:

> *Ontologically*, our position is identifiable as a form of *monism*, although, in view of the many complex varieties of ontology, the differentiation of its monistic from its pluralistic aspects is hardly worth the effort. If it is monism, it is *substantival* monism that we are talking about; yet it is *neutral*, and like Spinoza, we are prepared to apply *attributive pluralism* to the substance whenever our purposes might be served thereby. (p. 17)

One detects here a kind of cavalier nonchalance toward philosophical niceties. We assume, however, his view reflects the "neutral monism" advocated by William James (see Russell, 1961). James rejected the subject/object relation assumed to be fundamental in traditional philosophy, and rather proposed that the stuff the world is constructed of is neither mind nor matter, but rather something anterior to both. What might that anterior something be?

[1]These authors consider the kinds of extraction and compilation mechanisms that might account for the stabilities and variances in human perception and cognition, the kinds of stimulus transformations, if any, effected by our human hardware (e.g., neurology) and software (e.g., cultural programming). Probably for Kelly our human tendency to represent experience in contrasting categories was mainly a function of our neurological hardware, while culture played a large role in determining the content and ordering of those categories. In the midst of a dominant, positivistic S/R paradigm, he stood relatively isolated and proposed that we relied on abstract representations of stimuli, and constructions based on subjective goodness-of-fit criteria. Now others are attempting to model appropriate and explicit abstracting and compiling mechanisms, are attempting to populate the large "excluded middle" between stimulus bound and culturally determined perceptions and cognitions.

In reviewing Kelly's book Bruner (1956) says: "In a deep sense, the book reflects the climate of a generation of nominalistic thinking in the philosophy of science" (p. 355). Nominalism proposes that there is nothing universal except names, that is how we manage the anarchy of particulars. In Kelly's model, personal views of reality are constructed of the stuff of representations and constructs, the product of an interaction of perceiver and perceived. Nevertheless, while we rely on indirect representations of it, for Kelly (1955):

> The universe is real; it is happening all the time; it is integral; and it is open to piecemeal interpretation. Different men construe it in different ways. Since it owes no prior allegiance to any one man's construction system, it is always open to reconstruction. Some of the alternative ways of construing are better adapted to man's purposes than are others. Thus, man comes to understand his world through an infinite series of successive approximation. (p. 43)

In terms of a theory of science, we place Kelly with Popper, Campbell, and Simon, who believe that at least some of our knowledge is more than arbitrary tribal fabrication. Nevertheless, like Kelly, they criticize much of traditional epistemology, such as rational-empiricism and the inductive marshalling of supposedly objective positive instances. "Induction is scandalous, in that the psychological certainty of the conclusions we draw far outreaches the logical compellingness of the evidence" (Campbell, 1977, p. 15). But it is not induction per se that warrants attack, for by definition, going from the particular to the general involves going beyond the data, which all of us must do every minute of every day. Rather, it is the cavalier and uncritical reliance on general inductive inferences, in science, that is "scandalous." As Kelly (1955) notes, individuals navigate from the particular to the general (e.g., from the present to the future) by relying upon subjective, not objective, criteria in judging goodness-of-fit between our anticipations (inherent or acquired) and feedback.

Popper (1974a, 1974b) has led the way in criticizing the assumptions and practices of naive realism and traditional science, although retaining the firm conviction that some facets of "reality" can be mapped veridically if we radically revise our epistemology by, for example, eschewing naive induction and adopting his "falsifiability" criteria, which involve the conscientious search for negative evidence concerning your own hypothesis. Radical constructivists like Kuhn could point, however, to Kelly's Constructive Alternativism as providing, for the individual, a mode of construing that implies that much (most?) of our knowledge is constructed and is relative. Such knowledge is constructed to meet subjective personal and tribal expectations. For Kelly all our feedback consists of representations of events. In this sense he is a constructivist. But the model does not rule out various degrees of correspondence between our representations of events, on the one hand, and event structure and rhythm, on the other. Occasionally, even shadows deliver a lot of reliable and valid information. Kelly proposes that individuals share a common belief that "reality" can be grasped and, in our view, his theory does not preclude that possibility.

No conclusions without premises

The outlines of an emerging theory of knowledge suggests that much of our current knowledge is not only constructed but also "invalid," or of indeterminate validity in the traditional sense of external or construct validity. Much of our so-called knowledge consists of myths and small-world constructions, and is so context-specific as to have little generalizability. For the most part reliable knowledge derives from artificial structures of our own making in highly bounded laboratory domains, or symbol space (theories or models), which may, at best, reflect some degree of internal validity.

Heretofore, our models of knowledge have in the main focused on the explicit rational/empirical manipulations we apply to already reduced and simplified laboratory or questionnaire domains. Then, unwittingly or uncritically, we assume that such artificial or small worlds "represent" REALITY writ large. Understandably, working within these simplified domains, we developed appropriately simple models of knowledge based on "objective" observations combined by means of presumably unassailable logic into law-like statements. Such models were generated by logical positivism and by rational/empirical revisions in which absolute notions of truth were replaced by probabilistic ones. But in recent years positivism and its revisions face increasing criticism from a host of authors (Popper, 1970 & 1974b; Kuhn, 1970; Campbell, 1986; Brown, 1969; Merrell, 1982; Quine, 1969; Simon, 1983; Stockman, 1983,

etc.) whose works reflect a radically revised model of knowledge construction and revision.

The emerging theme is that our hypotheses, our observations and our logic are all based on a host of unexamined and unexaminable presuppositions, which Kelly calls core anticipations. In brief our knowledge is relative, not only to the *explicit* axioms, inference rules, and data samples employed, but also and more important to any array of implicit assumptions to which the explicit ones are anchored. Such implicit presuppositions typically lie beyond our conceptual reach (Simon 1983):

> the fallibility of reasoning is guaranteed both by the impossibility of generating unassailable propositions from particular facts, and by the tentative and theory-infected character of the facts themselves. Second, the principle of "no conclusions without premises" puts forever beyond reach normative statements. (p. 6).

The principle of "*no conclusions without premises*" — implicit as well as explicit premises — echoes Kelly's notion of no conclusions without anticipations and subjectively abstracted feedback, and is the underlying theme of this paper. Just as *explicit* axioms (hypotheses) and inference rules restrict the questions and conclusions that can legitimately be raised in formal debate or analysis (e.g., as in mathematics, logic, and laboratory studies), so too do *implicit* axioms (e.g., reductionist vs. holistic ontologies), and implicit inference rules (e.g. induction of positive instances vs. Popper's falsifiability principle) restrict the questions and conclusions that can be conceived of by those holding them, or in their hold. Our ontology (assumptions about reality), and our epistemology (assumptions about how to reliably map reality) strongly influence the kind of "reality" and "knowledge" we construct (Brown, 1969; Harre, 1985).

To address this emerging view of knowledge we require a model of construing that makes explicit provision for the role played not only by explicit assumptions, as in rational empiricism, but also for the very powerful influence that implicit assumptions, and subjective mechanisms exert in abstracting and combining selected fragments of evidence. Thus implicit assumptions shrink the "search and solution space" by focusing our attention on a drastically restricted range of questions, data, data transformation rituals, and possible solutions. We believe Kelly's personal construct theory (1955), in conjunction with modern decision theory, provides a useful theoretical foundation.

Kelly in a Modern Context

Kelly's work (1955) predated the emerging view that knowledge is relative. He saw knowledge as resulting from the compatibility, or the "goodness-of-fit," between our presuppositions and our abstracted and compiled representations or fragments of events. Truth for Kelly (1955) was subjectively construed: "Validation represents the compatibility (subjectively construed) between one's predictions and the outcome he observes. Invalidation represents incompatibility (subjectively construed) between one's predictions and the outcome he observes" (p. 158).

What are the prospects of such a subjective construing process producing valid knowledge? While many scholars echo Simon's theme that absolute validation or "certainty" about our models (constructions) lies beyond our grasp (Kuhn, 1970; Popper, 1974; Miller, 1978; Campbell, 1986), they disagree about how far beyond our grasp such certainty lies. These differences of opinion range from Popper's optimism (1970), through Kuhn's relativism (1970), to Collins' nihilism (1981a, 1981b). Also, as Kelly did before them, these scholars acknowledge the significant role played by anticipations and preconceptions in building and maintaining knowledge, although they disagree over the ratio of subjectivity to objectivity that forms the foundations of our "knowledge."

For example, while Popper was optimistic that we could draw more veridical maps of reality by replacing inductive inference with his falsifiability principle, he nevertheless echoes Kelly concerning our heavy reliance on "anticipations" when he notes that:

> Like Bacon, we might describe our own contemporary science — "the method of reasoning which men now ordinarily apply to nature" as consisting of "anticipations, rash and premature" and as "prejudices." (Popper, 1959, 278-279)

However, unlike Kuhn and Kelly, Popper views our rash and premature anticipations or prejudices as readily open to revision in the face of negative evidence, hence his plea for introducing his falsifiability principle in place of the current scandal of inducing positive instances.

Kelly's viewpoint, supplemented by those of some current investigators, provides helpful

guidelines for approaching the subjectivity/objectivity debate in both "small" or artificially closed systems, as well as in "large" open ones. In general, small worlds refer to artificial or highly controlled domains (e.g., technical production or laboratory studies), whereas large worlds refer to extensive space/time frames with many interacting, uncontrolled variables, and with multi-dimensional, delayed and/or fuzzy feedback. As noted earlier, for Kelly the foundations of human thought and knowledge rest not on facts but on subjectively construed goodness-of-fit between *anticipations* and *abstracted* representations of events.

> Much as he might like to shift responsibility, he [the theoretician] cannot ever claim that his theory is dictated by the facts it seeks to explain. The facts tell him only when he is wrong — and not always that. And besides there simply are too many ways of explaining the same facts — including a lot of ways that haven't turned up yet — for any of us to claim privileged communication with either God or Nature. (Kelly 1967, p. 155).

Feedforward vs. Feedback Mechanisms. How might such anticipations and subsequent abstracted feedback interact to produce different kinds of knowledge? May some knowledge approximate "objective truth" as in Popper's anticipations, while other knowledge reflects socially constructed truth, or tribal myths and puzzle-solving of the observer as in Kuhn's anticipations? One answer, reflected in the new philosophy of science (e.g., Brown, 1977; Harre, 1985), is that presuppositions, in the form of implicit axioms and inferences, serve as *feedforward* mechanisms that automatically direct attention, determine priorities, and select and compile data that is compatible with anticipations.

The goodness-of-fit between the feedforward axioms, for example species anticipations, and selected feedback may be "realistic" because of evolutionary selection of the neurological anticipations to fit typical environments. A case in point could be the relatively reliable anticipation by humans and birds, for example, that transparency and penatrability are highly correlated — glass being a later exception not provided for in the evolutionary data base (Campbell, 1987a). But the fit between anticipations and feedback can also be high not because feedback fits or reflects physical reality but because it is ambiguous enough that it can be selected and shaped to fit our personal, cultural, and scientific expectations. Such is commonly the case in inductive inferencing where citizens and scientists select positive instances and reject negative ones, from the flow of fuzzy feedback.

The importance of feedforward mechanisms, reflected in Kelly's thinking, is developed by Adams-Webber and Mancuso (1983). Conversely, rational/empirical models of science neglect the powerful role played by tacit presuppositions in knowledge construction and revision (Agnew & Brown, 1986). These positivistic models implicitly and explicitly focus attention on feedback without adequately addressing the host of problems encountered when feedback is fuzzy (e.g., delayed, multidimensional, or conflicting), as is the case with most problem solving and construction of non-trivial knowledge.

Feedforward mechanisms and Artificial Intelligence (AI)

Within cognitive science we see increasing attempts to model human problem solving by Artificial Intelligence programs and Expert Systems. The goal is to identify the ways knowledge is represented in intelligent systems — including animals, people or machines — and to transfer the extracted and compiled "intelligence" into computer based simulations and applications. Ambitious examples include the Japanese 5th Generation computer project and the American military, both devoting vast sums to basic and applied research. Some military and political leaders envision AI as transforming warfare into a kind of global video game involving robots (immune to bacterial and chemical warfare), computerized tanks, aircraft, and satellites, as in the United States Strategic Defence Initiative known as "Star Wars."

But for good or ill, AI to date fails to make provision for the potent role played by non-rational mechanisms in complex information processing (Agnew, Brown, & Lynch, 1986). As noted by Kelly (1967):

> Neither abstraction nor generalization has ever been computerized, nor can either be realized by an unimaginative obedience to the cannon of rationality, or by performing symbolic transformations of mathematics, useful as these procedures may otherwise be... the contribution the computer makes is to the economy of the language employed, not to the conceptualization... but housecleaning is not abstraction, and economizing does not constitute theoretical thinking. Ockam's razor is a surgical instrument, not a creative tool. (p. 155)

As a result, as noted by Hayes (1979) twelve years later:

> Artifical Intelligence is full of "toy problems": small artificial axiomatizations or puzzles designed to exercise the talents of various problem-solving programs or representational languages or systems. The subject badly needs some non-toy worlds to experiment with. (p. 242).

Not surprisingly "experts" within this exploding field differ on moral as well as theoretical and technical issues. The "neats" (e.g., see *AI Expert*, an applied journal) bound the field so as as to discount, for the moment, moral and non-rational issues, and so run the risk of being accused of working on toy problems. The "scruffies," on the other hand, cut a wider swath, exploring the role powerful, non-rational feedforward mechanisms play in human expertise, and experiment with the application of fuzzy, non-traditional logics (e.g., Minsky, 1983; Gaines, 1984). The neat/scruffie distinction, while highlighting the issue, excludes the large middle consisting of the majority of AI workers. In fact the simple distinctions involved in two-valued logic underlies most expert systems, which suggests that experts are people who radically reduce the search space, validly or otherwise.

Nevertheless, while expert systems were initially oversold, knowledge engineers are now including functional anticipatory axioms to identify and respond differentially to patterns in the feedback (e.g., Rennels & Shortliffe, 1987), whether those patterns reflect "natural" or socially contrived structures, and whether the editing be public (e.g., symptom groupings) or private (prognosis generation), witting (to fit an explicit theory) or unwitting (to fit professional bias). To what degree AI will be able to replace human skills (e.g., robotic soldiers), or supplement them (e.g., aiding physicians in diagnosis, treatment planning and evaluation), remains anyone's guess. Simon (1988) suggests that your guess will be based more on whether you are optimistically or pessimistically inclined (non-rational feedforward mechanisms?) than on hard data. If optimistic, as he is, you may see AI from a Promethean perspective, and believe that our burgeoning knowledge will help make us healthy, wealthy and wise. If pessimistic, you may entertain bleak Pandoran anticipations of moral and technical monstrosities slithering out of the AI box. At this stage AI represents an excellent example of Kelly's theory — subjectively construed goodness-of-fit criteria operating on a fuzzy data base in the service of strong anticipations.

Robust Anticipations According to Kelly and Kuhn

In general, the major difference between the comparatively neat rational/empirical models of knowing, on the one hand, and emerging scruffy constructivist models, on the other, lies in the relative emphasis each places on feedback (e.g., data), on the one hand, and feedforward mechanisms (presuppositions), on the other. In the rational/empirical model, feedforward mechanisms are seen to serve merely as working hypotheses, readily discarded or modified to reflect discrepant feedback. In the emerging view, presuppositions represent robust hidden-hand editors. They are robust in the face of anomalous or discrepant feedback, and they are hidden-hand editors in the sense that the individual or sub-culture employing them, or under their influence, remains unaware of the powerful role such preconceptions play in selecting supporting data, and discounting discrepant data. Such implicit editors not only help predetermine (construct) the knowledge we "discover," but also help to maintain it, to fill in the gaps, and to defend it (see Minsky, 1983; Piaget, 1985; Pascual-Leone, 1988 for discussion of powerful defence mechanisms designed to protect our hard-won and integrated cognitive structures).

As a clinician Kelly sought a model for understanding how individuals made sense not merely of the past and the current flood of ambiguous and conflicting experience, but of the future as well. Notice that the future is necessarily navigated via anticipations. Furthermore, in constructing models, it is primarily knowledge of the future we seek. Thus models of knowledge should focus attention especially on the design and maintenance of functional, and reality-oriented feedforward mechanisms, whether that reality is physically anchored and/or socially fabricated.

With regard to models of knowing Kelly's analogy is that of the "personal scientist" (Kelly, 1955; Shaw, 1980) who tests his or her anticipations against *abstracted representations* from personal experience. But Kelly is a clinician, so his "personal scientist" possesses powerful non-rational as well as rational problem solving

strategies. He used the concept of "core constructs" to identify strong anticipations (implicit axioms) that served as robust hidden-hand editors of individual experience. Here we draw the analogy with Kuhnian paradigms, which play a similar hidden-hand editing function for scientific schools. Kelly the clinician generates a psychological model, providing for bth rational and non-rational mechanisms, which helps anchor, at the individual level of analysis, Kuhn's historical theory of knowledge construction aimed at the cultural level of analysis.

We propose that a major distinction between the rational empiricists, on the one hand, and the constructivists, on the other, resides in different presuppositions (prejudices?) concerning the robustness of feedforward mechanisms. Kuhn, like Kelly, accepts the resiliency of our presuppositions (core constructs, beliefs, paradigms), while Popper, siding with the rational/empiricists on this particular issue, does not. Kelly points to the awesome robustness or resiliency of core constructs, noting that mode of construing can insulate us from experience, leaving us "gifted with a vast and versatile ignorance," or with "only one year of experience — repeated thirteen times" (Kelly, 1955, p. 171). Such core constructs, regardless of the external validity, can be functional for individual navigation, just as core paradigms can be functional for science, regardless of future judgements concerning their lack of external validity.

Kelly (1955) provides an example of a core construct that may well inform the construing of clients, therapists, and scientists alike:

> Our theme is the personal adventure of the men we are and live with — the efforts, the enterprises, the ontology of individuals so convinced there is something out there, really and truly, that they will not relent, no matter what befalls them, until they have seized it with their own hands. (p. 183)

Is this perhaps our most robust core construct, a powerful feedforward mechanism that helps humans in all cultures ignore or at least discount a history of personal, tribal, or scientific misperceptions, and to continue "reading" the shadows on the wall of Plato's cave as "reality seized with our own hands"? We see room in Kelly's "constructive alternativism" not only for functional and dysfunctional personal and tribal myths, but also for representations of knowledge that approximate some aspects of reality writ large. While the model is based on a subjectively construed fit between anticipations and abstracted feedback, and while much of our knowledge accordingly will be constructed or invented, some of our anticipations and abstractions may, by good luck, good rational/empirical management, or good evolution, be approximately "right." Abstracted feedback may provide isomorphic readings of "nature," rather than merely accurate reflections of tribal myths or personal prejudices. In brief, we see ample room in Kelly's formulation for both realists and constructivists. True, realists may not see nearly enough accommodation for themselves, and far too much for constructivists.

While Kelly's "alternative constructivism" suggests that feedback consists of abstractions in the service of anticipations, he readily acknowledges that not all such feedback is accommodating, and that, even subjectively construed, the goodness-of-fit can be bad. He was particularly concerned when a client's core constructs or anticipations were challenged, for here he saw the overwhelmed client "claw frantically for his basic construct" (Kelly, 1955, p. 167). Thus, when core constructs of the individual or paradigms for the scientist are "breached," they lose the hidden hand editing power of implicit feedforward mechanisms that serve to reduce and rationalize what would otherwise be overwhelming feedback of fragmented and conflicting experience. For Kelly, salient and/or massed negative feedback is the source of individual anxiety, and for Kuhn, the source of scientific uncertainty and paradigm revolutions. Also, in relation to the "frame problem," without such hidden hand editing, bits of negative evidence, for example a rejected manuscript, can resonate destructively throughout our system of beliefs, creating uncertainty and chaos. Alternatively when core constructs and paradigms are solidly in place, as they normally are, we unwittingly select positive instances and manage negative instances by conscientious, or trivial, rearrangements in our protective belt of auxiliary constructs or hypotheses (Lakatos, 1968) and by various damage-control mechanisms (Piaget, 1985; Pasqual-Leon 1988).

Infrastructure for Kelly's theory

Kelly's central concept of construing based on subjectively matching abstracted representations of events against anticipations and assumptions is reflected in subsequent literature: "A universe comes into being when a space is severed and

taken apart" (Brown, 1969, p. V); "Our computational scenarios, for all their quantitative detail, are only computing the consequences of the assumptions we have made according to the rules we have built in" (Gaines, 1984, p. 92); "The represention of an object as a collection of features is viewed as a product of a prior process of extraction and compilation" (Tversky, 1977, p. 329); "A person's processes are psychologically channelized by the way he anticipates events" (Kelly 1955, p. 46). For Kelly, human behaviour is primarily anticipatory rather than reactive.

Thus, by combining Kelly's concepts of anticipation, and of subjectively construed "fit," with the key concepts of current scholars, we elaborate the central hypothesis that to a significant degree we construct reality (knowledge). We do so, in keeping with our bounded rationality, by working in reduced search spaces using constructs with a finite range of convenience, constructs which enable us to functionally compute and justify the consequences of most of our assumptions and anticipations. By pooling the insights of Kelly the clinical psychologist, with those of Kuhn the scientific historian, Popper the philosopher, and Simon the decision theorist, we have the foundations for a general model of knowing. In this model, covering both individuals and scientific schools of thought, potent hidden hand editors drastically shrink complex problems to mind size. Within such reduced space/time frames we can then apply our limited rational/empirical tools for fine tuning the residual variance. On such simplified domains our bounded rationality can operate functionally, if not validly.

Kelly's theory, while brilliant, provides only a broad brush approach to the previously neglected question of how we shrink and simplify complex problem domains to mind size. Another critical question must be addressed. If our knowledge relies on robust feedforward mechanisms, and highly selective and abstracted feedback, then much of such knowledge must be highly fallible. How can we propose a general model of knowing with such an inherently flawed design, a design that is likely to produce invalid information about domain structure, a design that relies on a highly bounded and impoverished data base, consisting of abstracted fragments of feedback in the subjective service of a tangled or unspecified hierarchy of tacit anticipations and presuppositions?

We need to extend the model beyond Kelly's general postulates and clinical data base, we need to build a supporting conceptual infrastructure. For example: 1) We must specify what we mean by knowledge constructed in small- as opposed to large-world domains; (2) We must specify how confidence could be generated and maintained in constructed knowledge; 3) We must examine the conditions under which highly fallible knowledge can at the same time be highly functional; and 4) We must carefully examine the case for "realism" in relation to constructivism. We address the first two questions in the next two sections of this paper. We address the last two questions in a companion paper (Agnew & Brown, 1989).

Extracting and Compiling Knowledge in Small and Large Worlds

First we consider the kinds of small- and large-world data bases we encounter. For illustrative purposes, starting in an artificial or small world domain, we conceive of a simple knowledge structure as a relationship between two selected variables, X and Y, with some abstracting of positive and negative instances, or domain representations, serving to validate or invalidate our anticipations according to Kelly's subjectively construed goodness-of-fit criteria.

We can assume the following "if-then" decision rule for a class of binary choices where the possible outcomes are indicated in Figure 1:
If $x \geq Xc$ then accept (e.g., make decision A or choose action A)
if $x < Xc$ then reject (e.g., make decision B or choose action B)
If $y \geq Yc$ then conclude success (e.g., code as validating assumed relationship)
If $y < Yc$ then conclude failure (e.g., code as invalidating assumed relationship)

For example, assume that universities used an experimental paradigm to select graduate students, and evaluated the selection criteria in terms of various criteria of student performance. If there is a strong stable relationship (natural or constructed) between X (decision criteria) and Y (evaluation criteria) we would expect a large number of positive hits (students who pass selection criteria and succeed falling in cell 2) and negative hits (students who fail selection criteria and do not succeed falling in cell 3); while expecting relatively few false positives (students who pass selection criteria and do not succeed falling in cell 4); or false negatives (students

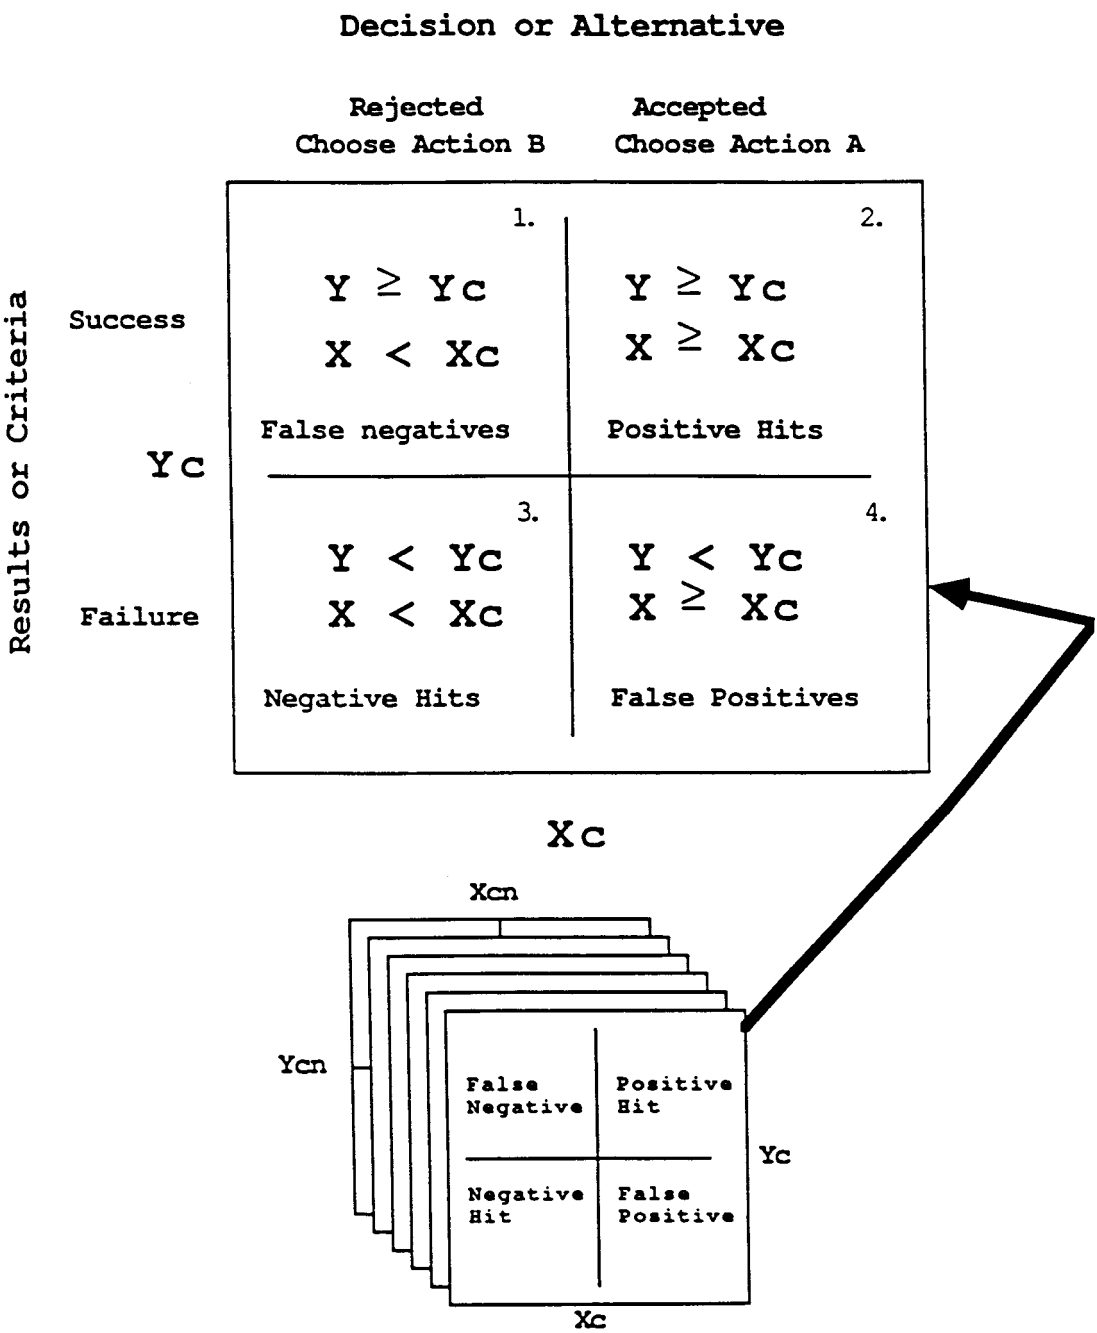

FIGURE 1

Two by two matrices representing the potentially infinite population of relationships between the X and Y constructs — populations from which specific samples are abstracted and compiled into "knowledge."

who fail selection criteria and succeed falling in cell 1). But, of course, in most non-trivial, large-world decision or action situations we don't have access to cell 1 or cell 3 data — data concerning rejected alternatives (Brown & Agnew, 1987). We rarely know what happens to students who do not meet the decision criteria and so are excluded from the data base. Furthermore the selection and evaluation criteria are typically multiple-dimensional and fuzzy and so readily manipulable and open to subjective goodness-of-fit criteria, as Kelly proposed. Yet to obtain valid estimates of a relationship between X and Y rationally requires reliable empirical feedback from data falling or edited into all four cells, as reflected in the traditional random assignment control group designs (Einhorn & Hogarth 1978).

Such symmetrical four-cell conditions would presumably produce "objective truth" in the traditional sense. But notice that to meet the conditions requires a radical restriction of the search space and of the number and type of variables to be considered — a costly trade-off of external for small world internal validity. Accordingly, rational/empirical truths are esentially limited to laboratory or small-world domains where artificial severances, extractions and compilations are conducted under highly selective conditions. But in what sense are such truths internally valid?

Any experiment typically represents a largely arbitrary selection of boundary conditions drawn from an unspecified population representing potentially infinite permutations and combinations of "representations" of the dependent and independent constructs (Yc... Yn, and Xc... to Xn, in Figure 1). As noted by Howard concerning basic research in perception (1974), if one is to avoid possible confounding, experimenters must be able to completely specify (logically and operationally) the search and solution space — in other words:

> One cannot have an idealized perceiver for an object in the world, for no investigator can be said to know all there is to know about a natural object. A theory of perception maps the descriptive structures of a perceiver not into the world but into an idealized description of some aspect of the world that the investigator creates. (p. 500).

Failure to appreciate this constructivist principle can give experimenters as false sense of confidence in their discovered/invented XY relationships. In reviewing the literature on orientation and shape perception Howard (1982) discusses various examples of unwitting confounding in otherwise carefully controlled research. Similarly in biological research it is becoming apparent (Sonea 1988) that the traditional method of studying bacteria in isolation, rather than in groups, drastically limits the validity of the results (e.g., relating to the development of resistant strains).

Review articles in the *Psychological Bulletin* typically document how varied, and confounded, various research studies of supposedly the same phenomenon turn out to be. In Howard's terms (1974, 1982) research studies represent illustrations of investigators unwittingly constructing and investigating different small worlds without appreciating the restrictive and confounding implications of their idealized descriptions. If more researchers acknowledged Howard's point that they are in effect constructing reality they would probably do a more thorough job of logically and operationally specifying the search and solution space, and so reduce the density of subsequently identified confounding, and reduce as well the necessary backtracking on even small-world generalizations. In brief, experimenters would be better hard-nosed empiricists if they recognized they were constructivists in Howard's sense. Tversky (1977) echoes Howard's point and summarizes the constructivist theme for perception in general: "The representation of an object as a collection of features is viewed as a product of a prior process of extraction and compilation" (p. 329). Such "prior processes" serve to reduce, radically, the space to be sensed or searched, and accordingly the space to be generalized to.

In brief, even in small two-by-two or highly controlled research domains, Kelly's constructive alternativism model appears to apply, with different experimenters unwittingly constructing different representations of "reality." Within such representation experimenters abstract and compile positive instances and so subjectively validate their anticipations. The key point is that even in basic research knowledge is extracted and compiled (constructed) in keeping with the tacit as well as the explicit anticipations and selective criteria of the investigators. Therefore, we see Kelly's model as applicable even in small and artificial worlds — the stronghold of traditional rational empiricism.

Of course, constructivist models have more obvious applications in larger and more open systems. In large worlds we rarely, if ever, have

access to data in all four cells of Figure 1, or the power to control and manipulate variables, and so fail to meet the normative rational/empirical conditions of validating or invalidating the predicted XY relationship. So not only must we rely on knowledge constructions based on tactic representations of reality, and on extracted and compiled fragments, as is the case in laboratory or artificial worlds, but we also lack access to feedback from two of the four cells — feedback necessary for "objective" (internal) validation of anticipated or hypothesized relationships. Briefly imagine applying the traditional rational/empirical model to the selection of graduate students, or the random assignment of graduate students to supervisors, while maintaining a blind control, that is, not disclosisng the student's status on selection criteria. Notice, that in designing such a "controlled" study, how many arbitrary (non-rational) decisions must be made in selecting a tiny sample of operational definitions, from the infinite population of proxies serving as potential "stand-ins" for the dependent and independent "constructs" (X... Xn, Y... Yn, in Figure 1). For example, a subset of such decisions might include: a) which combination of selection criteria to employ (e.g. grade point average, intelligence, high school grade), and which variables to ignore or control (e.g., age, sex, first language, first degree, of the host suggested by intuition and by the usually voluminous literature); b) which variables to manipulate if it is not a simple correlation study (e.g., "degree" and/or type of supervision); c) which criterion measures to apply when and for how long; d) how to code and analyze the results; and e) when to stop analyzing the results, etc., etc.

Any resulting extracted "truth" is therefore highly relative, relative to (dependent upon) the particular intended, and unintended, boundary conditions and controls (restrictions) of the study. It is relative to the selected and often arbitrary time/space severances, representations, compilations, and especially relative to the "success/failure" criteria selected and employed by the particular hypotheses testers or knowledge builders. How then can we continue to construe, with relative confidence, large worlds offering such impoverished feedback? The knowledge we primarily seek is about large worlds, for which we have access to only a restrictive and selective data base, as is the case in the domains of theory evaluation, construct and external validity, and linkages between the past and the future.

Confidence in Constructed Knowledge

To supplement Kelly's general theory of constructive alternativism, in addition to considering its application in small and large world domains as we have just done, we require an infrastructure that provides for degree of confidence ("subjectively construed validation") in various knowledge structures, particularly those based on feedback from limited and degraded data or representations — e.g., impoverished feedback available from only half the cells (cell 2 and cell 4) of the traditional four-cell rational/empirical model of validation.

Einhorn and Hogarth (1978) propose, and Brown and Agnew (1987) elaborate, such an infrastructure under the assumptions that: a) cell 1 and 3 data (Figure 1) are typically either unavailable or ignored (avoided); and b) outcomes are typically coded as frequencies rather than as probabilities. They then propose that confidence (C) in a relationship — subjectively construed validity — is related to feedback (F): $C = f(F)$ and that (F) in turn is a function of the weighted difference between positive feedback (np) and negative feedback (nf): $F = B1(np) - B2(nf)$ were the coefficients B1 and B2 reflect the relative reinforcing values — benefits and costs — of cell 2 positive hits on the one hand and cell 4 false positives (negative feedback) on the other, and where $B1 + B2 = 1.0$ ($B1, B2 > 0$).

These authors propose that in many situations, since the positive hit rate (i.e., unconditional probability) is greater than .5, in order to yield a significant negative feedback ($F < 0$) the B2 coefficient must be greater than .5 (e.g., B2 must be greater than B1). Thus, to discredit highly anticipated XY relations negative feedback must generate detectable penalties or costs higher than the payoffs or benefits generated by positive feedback. Therefore, to invalidate or reduce confidence in an anticipated relationship the negative feedback from cell 4 must have a highly negative impact on the construer or researcher. The false positives must generate penetrating contrasts (Kelly, 1955), or high negative surprise values (Gaines & Shaw 1985). Popper's falsifiability model reflects the same concept, but represents a normative proposal which optimistically ignores or discounts issues of bounded rationality, as well as ignoring the operation of

robust feedforward and damage control mechanisms (Lakatos, 1968; Minsky, 1983; Piaget, 1985). Mahoney (1988b), like Popper, makes stringent demands upon construers. In referring to constructivist metatheory he says:

> No other family of modern theories asks its adherents to maintain such a degree of self-examining openness, to so painstakingly tolerate and harvest (rather than eliminate) ambiguity, or to so thoroughly question both the answers and the questions by which they inquire. It is not easy to be a constructivist. (p. 18).

Like Popper, Mahoney has a remarkable capacity to harvest ambiguity and painstakingly seek and find higher order invariances. It is certainly hard to be a Mahonian constructivist, but perhaps not so hard to be a Kellyian constructivist; as we read Kelly, being a constructivist typically involves "doing what comes naturally." Mahoney is addressing the very thorny issue of what is to be done when doing what comes naturally doesn't work (e.g., when intrusive ambiguities cannot be accommodated within the range of convenience of the construer's core constructs, and when the development or reorganization of constructs leads to personal anxiety or the scientific angst of paradigm revolutions).

In brief, compatible conceptual infrastructures are emerging to supplement Kelly's general theory of subjective validation of our personal, cultural and scientific anticipations or hypotheses. In a companion paper (Agnew & Brown, 1989), we discuss the conditions under which highly fallible knowledge can be highly functional, and how a case for critical or hypothetical realism can be made within a constructivist frame of reference.

Finally, we view Kelly's constructivist perspective of knowledge as reflecting both humbleness and awe in his portrayal of human attempts to decode nature's formidable and convoluted puzzles. For us Kelly's approach reflects a spirit similar to that of an earlier and illustrious student of "truth":

> I do not know what I may appear to the world, but to myself I seem to have been only a boy playing on the seashore, and diverting myself in now and then finding a smoother pebble or a prettier shell than ordinary, whilst the great ocean of truth lay all undiscovered before me. (Issac Newton, quoted in Brewster, 1965, p. 407)

RÉSUMÉ

Les vues traditionnelles du savoir sont mises au défi. Une perspective "constructiviste" emergente telle que proposée par George Kelly, un ingénieur devenu clinicien, suggéré que nous pouvons en grande partie construire la réalité. Dans son «alternative constructive», Kelly présume que nous vérifions nos hypothèses et croyances en nous servant de critères convenables et bons «construits subjectivement» et appliqués aux différences que nous percevons entre les anticipations et l'information en retour. Son modèle de conceptualisation est compatible avec celui qui ressort de l'histoire et de la philosophie de la science et de la psychologie cognitive. Néanmoins, les constructivistes doivent répondre à cette question embarrassante, soit : coment le savoir faillible, construit comme il l'est, à partir de representations abstraites et incomplètes d'objets et d'événements, peut-il captiver et garder notre confiance comme il le fait et, en outre, démontrer qu'il est hautement fonctionnel, comme c'et le cas?

References

Adams-Webber, J. (1988). Personal communication.

Adams-Webber, J. (1979). *Personal construct theory: concepts and applications*. New York: Wiley.

Adams-Webber, J. & Mancuso, J.C. (1983). *Applications of personal construct theory*. New York: Academic Press.

Agnew, N. McK., & Brown, John L. (1989). Foundations for a model of knowing: II. Fallible yet functional knowledge. *Canadian Psychology, 30*(2), 168-183.

Agnew, N. McK., & Brown, John L. (1986). Bounded rationality: Fallible decisions in unbounded problem space. *Behavioral Science, 31*, 148-161.

Agnew, N. Mck., Brown, J.L., & Lynch, J.G. (1986). Extending the reach of knowledge engineering. *Future Computing Systems, 1*, 115-141.

Brewster, D. (1965). Memoirs of the life, writings, and discoveries of Sir Isaac Newton (Vol. 2). New York: Johnson Reprint Corporation.

Brown, G.S. (1969). *The laws of form*. London: George Allen & Unwin.

Brown, H.I. (1977). *Perception, theory and committment: The new philosophy of science*. Chicago: University of Chicago Press.

Brown, J.L., & Agnew, N. Mck. (1987). Feedback bias from ignoring the outcome of rejected alternatives. *Behavioral Science, 32*, 34-41.

Bruner, J.S. (1956). You are your constructs. *Contemporary*

Psychology, 1, 355-356.
Campbell, D.T. (1977). Descriptive epistemology: Psychological, sociological, and evolutionary. *William James Lectures*, Harvard University.
Campbell, D.T. (1986). Science's social system of validity-enhancing collective belief change and the problems of social science. In "Fiske, D.W., & Shweder, R.A. (Eds.), *Metatheory in the social sciences: Pluralisms and subjectivities* (pp. 108-135). Chicago: University of Chicago Press.
Campbell, D.T. (1987a). Neurological embodiments of belief and the gaps in the fit of phenomena to noumena. In Abner Shimony & Debra Nails (Eds.), *Naturalistic epistemology* (pp. 165-192). Dordrecht, Holland & Boston, MA: D. Reidel Publishing.
Campbell, D.T. (1987b). Selection theory and the sociology of scientific validity. In W. Callebaut & R. Pinxten (Eds.) *Evolutionary epistemology A multiparadigm program* pp. 130-158. Dordrecht, Holland & Boston, MA: D. Reidel Publishing.
Collins, H.M. (1981a). Stages in the empirical programme of relativism. *Social Studies of Science, 11*, 3-10.
Collins, H.M. (1981b). Son of seven sexes: The social destruction of a physical phenomenon. *Social Studies of Science, 11*, 3362.
Einhorn, H.J., & Hogarth, R.M. (1978). Confidence in judgement: Persistence of the illusion of validity. *Psychological Review, 5*, 395-416.
Fodor, J.A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, Mass: MIT Press.
Ford, K.M. (1988). Personal communication, July 1988.
Ford, K.M. & CHANG, P. (1989). An approach to automated knowledge acquisition founded on Personal Construct Theory. In *Artificial Intelligence Rsearch* (pp. 1-34). M. Fishman (Ed.). Greenwich Conn.: J.A.I. Press.
Gaines, B.R. (1984). Methodology in the large: Modeling all there is. *Systems Research, 1*, 2, 91-103.
Gaines, B.R., & Shaw, M.L.G. (1984). Emotion: a systematic development of the personal construct model. Paper presented to the Personal Construct Network Conference, Cincinnati, June 22-24, 1984.
Harre, R. (1985). *Philosophies of science*. New York: Oxford University Press.
Howard, I.P. (1974). Proposals for the study of anomalous perceptual schemata. *Perception, 3*, 497-513.
Howard, I.P. (1982). Human visual orientation.Chichester, Sussex: John Wiley.
Jones, M.R. (1976). Time, our lost dimension: Toward a new theory of perception, attention, and memory. *Psychological Review, 83*, 323-355.
Kelly, G.A. (1955). *The Psychology of personal constructs*. Volume 1: A Theory of Personality, Volume 2: Clinical Diagnosis in Psychotherapy. New York: W.W. Norton.
Kelly, G.A. (1967). The role of classification in personality theory. In M.M. Katz, J.O. Cole & W.E. Barton (Eds.), *The role and methodology of classification in psychiatry and psychopathology* (pp. 155-164). Washington, D.C.: Government Printing Office.
Kuhn, T. (1970). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
Lakatos, I. (Ed.). (1968). *The problem of inductive logic*. Amsterdam: North Holland.
Mahoney, M.J. (1988a). Constructive metatheory: I. Basic features and historical foundations. *International Journal of Personal Construct Psychology, 1*, 1-35.
Mahoney, M.J. (1988b). Constructive metatheory: II. Implications for psychotherapy. Invited paper presented at the *Seventh International Congress on Personal Construct Psychology*, August 5-9, 1988, Memphis, Tennesse, Psychotherapy Section, R. Neimeyer (Chair).
Merrell, F. (1982). Semiotic foundations: *Steps toward an epistemology of written texts*. Bloomington: Indiana University Press.
Miller, J.G. (1978). *Living systems*. New York: McGraw-Hill.
Minsky, M. (1983). Jokes and the logic of the cognitive unconscious. In Groner, R., Groner, M., & Bischof, W.F. (Eds.), *Methods of heuristics* (pp. 171-193). London: Lawrence Erlbaum.
Pascual-Leone, J. (1988). Affirmations and negations, disturbances and contradictions, in understanding Piaget: Is his later theory causal? *Contemporary Psychology, 33*, 420-421.
Piaget, J. (1985). *The equilibrium of cognitive structures: The central problem of intellectual development*. Chicago: University of Chicago Press.
Popper, Karl. (1959). The logic of scientific discovery. New York: Basic Books.
Popper, Karl. (1970). Normal sciences and its dangers. In Lakatos, I and Musgrave, A., (Eds.), *Criticism and growth in knowledge* (pp. 49-58). Cambridge: Cambridge University Press.
Popper, K.R. (1974a). Autobiography of Karl Popper. In P. Schilpp, (Ed.), *The philosophy of Karl Popper, Book I* (pp. 2-181). LaSalle, Ill: Open Court.
Popper, K.R. (1974b). The philosopher replies. In P. Schilpp, (Ed.), *The philosophy of Karl Popper*, (pp. 961-1148). LaSalle, Ill: Open Court.
Pylyshyn, Z.W. (1987). *The robot's dilemma: The frame problem in artificial intelligence*. Norwood, N.J.: Ablex Publishing Corporation.
Quine, W.V. (1969). *Ontological relativity*, New York: Columbia University Press.
Rennels, G.D. & Shortliffe, E.H. (1987). Advanced computing for medicine. *Scientific American*, October, 154-161.
Russell, B. (1961). *History of western civilization*. London: George Allen and Unwin.
Rychlak, J.F. (1987). Following the Skinnerian example: In search of a science of psychology beyond the black-box computer. Paper presented at the *American Psychological Association Meeting* in New York in August 1987, Social Psychology Section, K. Gergen (Chair).
Sonea, S. (1988). The global organism. *The Sciences*, July/August, 38-45.
Shaw, M.L.G. (1980). *On becoming a personal scientist*. London: Academic Press.
Simon, H.A. (1977). Models of discovery. Boston: D. Reidel.
Simon, H.A. (1983). *Reason in human affairs*. Stanford: Stanford University Press.
Simon, H.A. (1988). Ideas. *Canadian Broadcasting Corporation*. Toronto, July 21, Artificial Intelligence, L. Sinclair (Host).
Shimony, A. (1981). Integral Epistemology. In M.B. Brewster & B.E. Collins (Eds.), *Scientific Inquiry and the social sciences: A volume in honour of Donald T. Campbell* (pp. 98-123). San Francisco: Jossey-Bass.
Tversky, A. (1977). Features of similarity. *Psychological Review, 84*: 327-352.

Von Wright, G.H. (1966). The paradoxes of confirmation. In J. Hintikaa & P. Suppes (Eds.), *Aspects of inductive logic* (pp. 208-218). Amsterdam: North Holland.

Weizmann, F. (1987). Personal communication, November, 1987, York University, Toronto. A file of Kelly's manuscripts including: Personal construct theory and the psychotheraputic interview, an undated manuscript from which this quote was taken.

Wimsatt, W.C. (1981). Robustness, Reliability, and overdetermination. In M.B. Brewster & B.E. Collins (Eds.), *Scientific Inquiry and the social sciences: A volume in honour of Donald T. Campbell* (p. 163). San Francisco: Jossey-Bass.