problems, for example, are genetically determined, as is only emphasized by Anderson's insistence on the shaping role of the environment.

Given that the optimizing nature of behavior is not immediately obvious, nor entailed by evolutionary theory, it makes most sense to consider some empirical evidence that sheds light on the issues. In fact, a large amount of evidence has been collected that indicates less than optimal behavior in many situations. Let us touch on the issues of mental effort, local versus global optimization, and problem specification, all of which figure prominently in the target article.

According to Anderson, one of the key predictors of human performance in a rational theory is the cost of mental effort. But numerous studies of behavior indicate that people are often willing to expend more mental effort on heuristics that are in fact less ideal for the task at hand. For example, Kahneman and Tversky (1984) show that people who are about to pay $15 for a calculator are more willing to drive to another store to buy the calculator for $5 less than is a second group of people who are about to buy a calculator for $125. Although the full account of this behavior is beyond our present purpose (for a discussion, see Kahneman & Tversky 1984; Shafir et al. 1989; Thaler 1980), suffice it to note that, in this and other problems of its kind, people deviate from the standard rational theory of consumer behavior by working harder to evaluate gains and losses in relative rather than absolute terms. By basing their decision on the same $5 difference rather than the ratio it forms of the total price, subjects would save on mental effort and conform with the rational model, and this is what they do not do. In a similar vein, when asked to estimate the likelihood that a totally uninformative description of a person is of an engineer, subjects use the more effortful representativeness heuristic, rather than simply relying on the base rates that are given and that would lead to a normatively more adequate response (Kahneman & Tversky 1973). Although in Anderson's rational analysis "cognitive performance maximizes the difference between the expected gain and cost of mental effort," subjects in the experiments above do not seem confined to such considerations.

Anderson argues that whereas local optimality is normally achieved, global optimality is sometimes foregone for reasons such as limited memory. But a number of studies show that the discrepancy between local and global strategies really is not so easy to resolve. Redelmeier and Tversky (1990; see also Slovic et al. 1978), for example, asked hundreds of clinicians to make treatment choices pertaining either to a single patient or to a group of comparable patients. The clinicians weighted certain criteria (such as personal concerns of the patient and cost effectiveness) differently in the two cases and, as a result, exhibited different preferences between treatments. They preferred one treatment when dealing (locally) with one patient, and the other treatment when contemplating (globally) the entire group. Although their local and global preferences are clearly discrepant, it is not at all clear that one is "right" and the other "wrong": It may just be that we give relatively more weight to the human dimension when patients are in our office, and less when we contemplate the finances of a public health policy. Although not licensed by Anderson's rational approach, conflicting local and global strategies may just be a natural outcome of the way we process information.

A precise description of the environment is a critical part of the rational theory. Work on "framing" has shown, however, that the same description of an environment leads to different behaviors when framed differently. Thus, when asked to choose between two alternative treatments of a disease, people prefer one treatment if the outcomes are framed in terms of lives lost, and the other treatment if the problem is framed in terms of lives saved (see Kahneman & Tversky 1984; McNeil et al. 1982). Such violations of "description-invariance" as well as of "procedure invariance" (where people express discrepant preferences depending on the particular elicitation procedure that is used; see Slovic et al., 1982, for a discussion) have now been documented in numerous domains, in hypothetical as well as real world situations, with both high and low stakes, and both with and without monetary incentives. These behaviors are not going to go away, and they do not seem to point us in the same direction as Anderson's conclusion "that many of the major characteristics of human cognition can be explained as an optimal response, in a Bayesian sense, to the informational structure in the world."

The rationality of causal inference

Thomas R. Shultz

Department of Psychology, McGill University, Montréal, Québec H3A 1B1, Canada

Electronic mail: *ints@musicb.mcgill.ca*

Rational analysis could be viewed as an attempt to rise above some of the major debates in contemporary cognitive science concerning the best mechanistic explanations of cognition. It does this by showing that a wide variety of cognitive phenomena are optimal responses to goal satisfaction in particular environments given certain minimal computational limitations.

The breadth of intended coverage for rational analysis is impressive. At the current rate of application, we may soon be asking whether it qualifies as a candidate for a unified theory of cognition. One reason the approach seems so general is that it operates at a more abstract level than the more mechanistic candidates for unified theories.

Although Anderson admits that ordinary humans do not explicitly undertake Bayesian computation, he claims that rational analysis can easily be converted into plausible reasoning mechanisms. It is doubtful, however, that the constraints supplied by rational analysis are sufficient to favor particular mechanisms. Much current debate centers around the fact that a fair number of distinct algorithms account for many of the phenomena listed by Anderson. Rational analysis will satisfy researchers used to mechanistic accounts only insofar as it supplies sufficient mechanistic constraints.

One of the four main areas of application in the target article is causal inference, an area that optimizes my own interest and expertise. Anderson specifies his analysis of causal inference in two concise Bayesian equations (9 and 10). As with many Bayesian analyses, it is difficult to understand how the probabilities on the right sides of these equations are any more fundamental than those on the left sides. For example, in Equation 10, the probability of a rule applying in the presence of cues seems no less fundamental than the probability of the cues if the rule did apply. What evidence would support Anderson's view that the reverse is true? What sort of reasoning mechanism would conform to this arrangement of conditional probabilities?

Anderson makes a very useful point when he argues that cues to causality depend on underlying causal models (cf. Doyle, 1989, for a more elaborated version of this). He applies this notion a bit too narrowly, however. For example, in Phenomenon 15, Anderson argues that similarity is used when the subject holds a causal model specifying that a cause transfers part of itself to the effect, as in coloring phenomena. But we have found that similarity is also used in contexts of number and size (Shultz & Ravinsky 1977). These cases too could be understood as using similarity only when it conforms to a causal model, but this would require a somewhat wider range of underlying causal models.

Anderson stresses that in a rational analysis the goal of causal inference is to predict future events. But causal inference traditionally involves explanation and planning as well as prediction. People may be interested in explaining an event by finding its cause even if the need for predicting the event never arises. Once causal relations are known, they can also be used in the construction of plans to satisfy goals. As Anderson (1990) realizes, this brings us into the related realm of problem solving. My point is that an unwarranted theoretical emphasis on prediction leads to the unfortunate theoretical neglect of explanation and planning.

Indeed, the only experiment Anderson reports here in much detail (Phenomenon 14, temporal and spatial contiguity) is one that involves explanation (or attribution, as it is more commonly known). Phenomena 13 (contingency) and 15 (similarity) also involve causal attribution rather than prediction. With respect to Phenomenon 16 (generalization), Anderson (1990) does present some new and interesting experiments on prediction, but it seems to me that the goals for causal inference ought to include explanation at least, and probably planning, as well.

Phenomenon 16, that a causal law generalizes to objects of the same category, seems obvious from many perspectives and so is not a unique prediction of a Bayesian analysis. It is not a well studied problem, and Anderson's efforts may stimulate more research.

Phenomenon 13, ignoring joint nonoccurence information in contingency analysis, could likewise be explained by people's tendency to avoid negative information, presumably because of its complexity (Schustack & Sternberg 1981), or to ignore nonevents, presumably because they cannot ordinarily be detected. Anderson's Bayesian explanation of this phenomenon appears to work less well than the foregoing explanations, first, because it has subjects assuming the causal relation they are trying to detect, and second, because it contradicts the well documented tendency of people to ignore prior probabilities (Tversky & Kahneman 1980). The tendency to ignore priors could bode ill for any account that postulates Bayesian mechanisms.

The fact that many of these phenomena have various mechanistic explanations is symptomatic of the fact that a rational analysis under-constrains mechanistic accounts. Although it may always be possible to design probability values that conform both to Bayesian formulas and to behavioral data, the exercise is not terribly satisfying unless it distinguishes a mechanistic account. If the mechanistic account turns out to involve computations on probabilities, then it behooves the researcher to use behavioral data to constrain the choice of probabilities more directly, rather than using probabilities that simply make the equation fit the data to be explained.

So far, Anderson provides no clear analysis of the time course of causal knowledge. Is novel behavior typically as optimal as well-practiced behavior? Eventually, this may get cashed out in terms of the progressive refinement of prior probabilities, but it would have been interesting to see this done on some well known developmental phenomena.

Anderson identifies all of his causal inference phenomena (13–16) as having previously been characterized as irrational. It is extremely useful to show how such phenomena can be alternatively described as rational under differing assumptions. Anderson could well have included research by Kuhn (1989) on the *errors* people make in using the covariation heuristic in causal reasoning.

But how does the brain think? (or wasn't that the question?)

Steven L. Small

Department of Neurology, University of Pittsburgh, Pittsburgh, PA 15261
Electronic mail: *sls@dsl.pitt.edu*

Introduction. Anderson describes a way to analyze cognition in terms of human behaviors that are optimized to particular features of the external environment. He refers to the method as "rational analysis"; I refer to the underlying principle as either the "adaptation" or "optimization" assumption. Using rational analysis, he investigates several aspects of human cognition. I believe that the formal statement of the adaptation assumption adds a useful perspective to our thinking about cognition, but that the formal application of the method in rational analysis does not explain the kinds of things that need explaining.

Adaptation and cognitive science. Human cognition is as likely a priori to be adaptive as other human brain processes, such as regulating blood pressure or seeing in color. That it can be shown to be adaptive is interesting, but does not lead to an understanding of how it works. Cognitive scientists have always operated under some form of the optimization assumption, yet they believed that their mission was to explain not how the behaviors were optimized, but how the mind actually produced the behaviors. Similarly, cognitive neuroscientists focus on the phylogenetic and ontogenetic development of the brain and how the resulting brain structures lead to cognitive behavior.

Definitions and theory application. Anderson begins by stating that the goal of cognitive science is to predict the "output of human cognition." After presenting his method of analysis of cognitive behavior, he applies it to four areas in cognition. A number of questions can be raised about the definition and the areas chosen for analysis.

First, cognitive behaviors are not exclusively "outputs"; in fact, many of the tasks studied by cognitive scientists, such as reading (Just & Carpenter 1980), listening to music, understanding visual scenes (Marr & Poggio 1977), and inventing new scientific theories (Langley et al. 1983), cannot readily be described as outputs. But there is no doubt that they are cognitive tasks.

Second, four aspects of cognition were chosen for an analysis according to the adaptive theory: (a) recall memory; (b) categorization; (c) causal inference; and (d) problem solving. These tasks have something in common that makes them easily analyzed in terms of optimization; they are (generally) conscious, symbolic (propositional), and controlled processes.

Third, the data best applied to the optimization theory also have idiosyncrasies in common: They come from carefully controlled studies of goal-directed, rewarded, rational, propositional behaviors. Anderson thus provides an interesting account of some statistical regularities of certain conscious goal-directed behaviors. Although this might be useful as a descriptive theory, it does not help in uncovering the explanatory theory sought by a majority of researchers.

By defining cognition in a narrow way, excluding automatic aspects of visual and linguistic processing, for example, some of these objections can be averted. It is undesirable, however, to restrict the notion of cognition in this way, as it excludes some of the best understood aspects of overall behaviors that are universally considered cognitive. To continue the previous examples, the understanding of spoken language or of visual scenes requires both so-called "low level" and "high level" cognitive abilities. The interdependence of such processing "levels" requires their coordinated consideration in the development of cognitive theories.

Neuroscience. Although some cognitive researchers are reluctant to include basic neurobiological concepts in the development of cognitive theories, there is increasing interest in using data from neurobiology to constrain theory construction in cognitive science. The two sides of this coin are (a) the use of basic neuroscientific data from researchers in anatomy, physiology, and pharmacology; and (b) the consideration of clinical neuroscientific data (from neurology and psychiatry) about people with obviously abnormal cognitive systems and specific neurobiological defects. Must cognitive scientists restrict themselves to the brain as a black box? The optimization assumption side-steps the entire issue of how the mind/brain works, but that is in fact the main issue for a large proportion of cognitive scientists.

In this regard, do maladaptive and pathological cognitive