In sum, whether Take-The-Best is fast, is architecture dependent, and whether frugality is a virtue are questioned by the many other cognitive functions that require fast, parallel, and integrative approaches. Moreover, the weak empirical basis for Take-The-Best seems consistent with people normally adopting an integrative approach (albeit with limited cues).

Sub-optimal reasons for rejecting optimality

David R. Shanks and David Lagnado

Department of Psychology, University College London, London WC1E 6BT United Kingdom. **{d.shanks; d.lagnado}@ucl.ac.uk**
www.psychol.ucl.ac.uk/david.shanks/Shanks.html.

Abstract: Although we welcome Gigerenzer, Todd, and the ABC Research Group's shift of emphasis from "coherence" to "correspondence" criteria, their rejection of optimality in human decision making is premature: In many situations, experts can achieve near-optimal performance. Moreover, this competence does not require implausible computing power. The models Gigerenzer et al. evaluate fail to account for many of the most robust properties of human decision making, including examples of optimality.

A paradox in the rationale for fast and frugal algorithms. There is a curious paradox in Gigerenzer et al.'s argument for the role of fast and frugal algorithms in human decision-making (Gigerenzer et al. 1999). They suggest that psychologists have been led astray by focusing on behavior from an optimization perspective and they imply that optimization models are implausible, intractable, and require demonic capacities. Instead, they urge us to explore fast and frugal algorithms. Yet they admit on p. 237 that categorization performance can often be optimal, but if that is the case, then surely the case for fast and frugal algorithms – which will almost certainly never achieve this level of performance – evaporates?

To evaluate Gigerenzer et al.'s case for minimal complexity in cognitive processes it is critical to determine whether decision making is truly optimal. The jury is still out on this issue, of course, but what is indisputable is that near-optimal performance can be achieved by experts in many realms including categorization (Anderson 1991; Ashby & Maddox 1992) and choice (Binmore 1991; Davis et al. 1993).

Gigerenzer et al. repeatedly ridicule what they take to be "optimal" theories (e.g., multiple linear regression, MLR) on the grounds that they require unrealistic amounts of computation (e.g., p. 76), but this is a highly misleading claim. Contrary to the impression made by Gigerenzer et al., it is possible to find a regression solution in minimal time without doing any computation at all. Imagine a set of points each represented by a peg on a two-dimensional board (the solution also works in principle in *n* dimensions). Then attach a long thin rod by a set of elastic bands to the pegs. By minimizing the allocation of tension across the elastic bands, the rod will align exactly according to the regression equation. As another example, consider the well-known Travelling Salesman problem in which the shortest route must be computed that visits each of a number (*N*) of cities exactly once. Despite its computational complexity (the computing time needed to solve this problem increases faster than any power of *N*), near-optimal solutions can be achieved by parallel neural networks in the blink of an eye (Hopfield & Tank 1985). What objection is there to the view that the human cognitive system approximates optimality by use of parallel constraint-satisfaction processes?

It is also troubling that Gigerenzer et al. take multiple linear regression (MLR) as one of their benchmark models throughout the book. Humans can learn highly nonlinear judgment rules in a variety of domains (Ashby & Maddox 1992; Ceci & Liker 1986) so MLR is simply not an appropriate model. If TTB (Take The Best) and CBE (categorization by elimination) approximately match the performance of MLR and if human experts significantly outperform MLR then the obvious conclusion is that TTB and CBE are inadequate models of human performance.

Implausibility of the CBE model. We believe that the candidate fast-and-frugal model for categorization which Gigerenzer et al. present, the CBE model, is wholly inadequate for human performance. First, it is unable to predict one of the benchmark phenomena of categorization, namely the ubiquitous "exemplar effect," that is, the finding that classification of a test item is affected by its similarity to specific study items, all else held constant (e.g., Whittlesea 1987). Even in the case of medical diagnosis, decision-making in situations very like the heart-attack problem Gigerenzer et al. describe is known to be strongly influenced by memory for specific prior cases (Brooks et al. 1991). The recognition heuristic is not adequate to explain this effect because it concerns the relative similarity of previous cases, not the absolute presence versus absence of a previous case. If a heavy involvement of memory in simple decision tasks seems to characterize human performance accurately, then plainly models which ignore this feature must be inadequate.

Second, there is strong evidence against deterministic response rules of the sort embodied in Gigerenzer et al.'s fast-and-frugal algorithms (Friedman & Massaro 1998): for instance, Kalish and Kruschke (1997) found that such rules were rarely used even in a one-dimensional classification problem. Thirdly, CBE is not a model of learning: it says nothing about how cue validities and response assignments are learned. When compared with other current models of categorization such as exemplar, connectionist, and decision-bound models, which suffer none of these drawbacks, the CBE model begins to look seriously inadequate.

Methodology of testing the models. By taking a tiny domain of application (and one which is artificial and highly constrained), Gigerenzer et al. find that the CBE model performs competently and conclude that much of categorization is based on the application of such algorithms. Yet they mostly do not actually fit the model to human data. The data in Tables 5-4, 11-1, and so on, are for objective classifications, not actual human behavior. It is hard to see how a model's ability to classify objects appropriately according to an objective standard provides any evidence that humans classify in the same way as the model.

Even in the cases they describe, the models often seriously underperform other models such as a neural network (Table 11-1). Compared to the more standard approach in this field, in which researchers fit large sets of data and obtain log-likelihood measures of fit, the analyses in Chapters 5 and 11 are very rudimentary. Gigerenzer et al. report percent correct data, which is known to be a very poor measure of model performance, and use very small datasets, which are certain to be highly nondiscriminating. The difference between the CBE model and a neural network (e.g., up to 9% in Table 11.1) is vast by the standards of categorization research: for instance, Nosofsky (1987) was able to distinguish to a statistically-significant degree between two models which differed by 1% in their percentages of correct choices.

Melioration as a fast-and-frugal mechanism. The algorithms explored by Gigerenzer et al. (TTB, CBE, etc.) share the common feature that when a cue is selected and that cue discriminates between the choice alternatives, a response is emitted which depends solely on the value of that cue. Gigerenzer et al. (Ch. 15, p. 327, Goodie et al.) consider the application of such models to the simplest possible choice situation in which a repeated choice is made between two alternatives in an unchanging context. The prototypical version of such a situation is an animal operant choice task in which, say, a food reinforcer is delivered according to one schedule for left-lever responses and according to an independent schedule for right-lever responses. As Gigerenzer et al. point out (p. 343), fast-and-frugal algorithms predict choice of the alternative with the highest value or reinforcement rate. Although this may seem like a sensible prediction, human choice does not conform to such a "momentary-maximization" or "melioration" process. In situations in which such a myopic process does not maximize overall reinforcement rate, people are quite capable of adopting better response strategies. A well-known example is the Harvard Game (Herrnstein 1997) in which one response alternative (say, right) always pays more at any moment than the other (left), but where overall reinforcement rate is maximized by allocating all responses to left. People's behavior is often seen to approach optimality under such conditions (Herrnstein et al. 1993; R. Tunney & D. Shanks, unpublished data). Yet again we have an example of humans' ability to achieve near optimal levels of performance, exactly as the "demonic" theory of rational choice predicts.

The selection problem. Gigerenzer et al. say very little about how individual heuristics are selected for application to specific problem domains. Such meta-level decisions will typically require some prior knowledge about the structure of the environment (e.g., whether it is non-compensatory, J-shaped, etc.), which may add substantially to the overall processing costs of a fast and frugal model. This would reduce its advantage over those models that can learn about the environment and have general applicability (thus cutting out the metadecision stage). In the majority of the simulations in the book, one particular heuristic is pre-selected to operate in a specific environment. More rigorous tests would place a complete fast and frugal system (different heuristics plus metadecision heuristics) in a variety of different environments, and compare its performance against an alternative general-purpose learning model.

The precision/accuracy trade-off. On a more positive note, we welcome Gigerenzer et al.'s shift of emphasis from "coherence" to "correspondence" criteria. This is an important step towards a more complete understanding of rationality, and removes some of the obstacles placed by the heuristics-and-biases school. In addition to the examples cited in the book, the inadequacy of coherence criteria has been demonstrated in various experiments in which people trade the probability of being correct for increased precision in their judgments (Yaniv & Foster 1995), thereby sacrificing probabilistic coherence for a possible gain in informational content. It is not clear, however, how readily this fits into the fast and frugal picture. Computing a trade-off between precision and accuracy would appear to place additional processing demands on the decision maker, contrary to the spirit of speed and frugality.

We believe that this problem is resolvable by identifying the appropriate correspondence criterion, and the cognitive mechanisms attuned to this criterion. Recent work in human causal induction (López et al. 1998) suggests that predictive judgments are mediated by associative mechanisms sensitive to real-world statistical contingencies. Furthermore, it can be shown that networks sensitive to such a measure automatically compute a precision/accuracy trade-off. This suggests, contra Gigerenzer et al., that although the mechanisms underlying our mental algorithms may be simple, the computations which they embody need not be.

Fast and frugal heuristics: What about unfriendly environments?

James Shanteau and Rickey P. Thomas

Department of Psychology, Kansas State University, Manhattan, KS 66506-5302 **{shanteau; thomas}@ksu.edu** **www.ksu.edu/isbr/shanteau.htm**

Abstract: *Simple heuristics that make us smart* offers an impressive compilation of work that demonstrates fast and frugal (one-reason) heuristics can be simple, adaptive, and accurate. However, many decision environments differ from those explored in the book. We conducted a Monte Carlo simulation that shows one-reason strategies are accurate in "friendly" environments, but less accurate in "unfriendly" environments characterized by negative cue intercorrelations, that is, tradeoffs.

Text removed due to third party copyright

Text removed due to third party copyright