There are several problems with this; they point to a difficulty with the research programme as a whole. First, there is a very important equivocation in how the stockmarket experiment is being described. The authors write as if their strategy was a pure application of the recognition heuristic. But this seems wrong. They did not invest in companies that they recognised. Rather, they invested in companies that had a high national and/or international recognition factor, where this is calculated statistically by comparing the recognition judgments of several different populations. These are two very different things. The first would have been a pure fast and frugal heuristic. The second, in contrast, seems much closer to a calculated investment strategy. What makes this equivocation important is that the notion of rationality applies very differently in the two cases. It is hard to see how anything other than a pure success-based criterion of rationality could be applied to the fast and frugal version of the recognition heuristic. Or, to put this another way, it is hard to see what reasons there might be for holding that it is rational to make one's investment decisions solely according to whether one has heard of the companies in question other than that the strategy more or less works over time – and it is equally hard to see how, if the strategy doesn't work, it could then possibly be described as rational. But the same does not hold of the sophisticated investment strategy of investing only in companies with a statistically attested high recognition. There are all sorts of reasons why this is a rational strategy to adopt – quite apart from the well-documented "big company effect" in bull markets (to which the authors themselves draw attention) and the simple thought that a company with a high recognition factor will correspondingly have a high market share. That is to say, even if it did turn out (as it probably would in a bear market) that the strategy did not beat the index it might well still count as a rational strategy to have adopted.

What this points us to is an important discussion in the concept of rationality. A workable concept of rationality must allow us to evaluate the rationality of an action without knowing its outcome. Without this the concept of rationality cannot be a useful tool in the control, regulation and evaluation of decision-making as and when it happens. And it is precisely such a way of evaluating the rationality of an action that we are offered by the orthodox normative theories of expected utility maximisation and so forth. But is far from clear that Gigerenzer and his co-workers have offered a genuine alternative to this. They claim to have replaced criteria of rationality based upon logic and probability theory with a heuristic-based criteria of real-world performance. but it doesn't look as if they've offered us criteria of *rationality*.

NOTE

1. For discussion of ways to strike the balance between these two facets of rationality see Bermúdez 1998; 1999a; 1999b; 1999c; 2000 and the essays in Bermúdez and Millar (in preparation).

How smart can simple heuristics be?

Nick Chater

Department of Psychology, University of Warwick, Coventry, CV4 7AL, United Kingdom **nick.chater@warwick.ac.uk**

Abstract: This commentary focuses on three issues raised by Gigerenzer, Todd, and the ABC Research Group (1999). First, I stress the need for further experimental evidence to determine which heuristics people use in cognitive judgment tasks. Second, I question the scope of cognitive models based on simple heuristics, arguing that many aspects of cognition are too sophisticated to be modeled in this way. Third, I note the complementary role that rational explanation can play to Gigenerenzer et al.'s "ecological" analysis of why heuristics succeed.

Gigerenzer, Todd, and the ABC Research Group have provided a series of impressive demonstrations of how simple "fast and frugal" cognitive heuristics can attain surprisingly impressive levels of performance, comparable to human performance in a range of tasks. They show, for example, that decision making based on a single piece of evidence, rather than integrating across all available evidence, can lead to close optimal performance in a wide range of estimation tasks (Gigerenzer et al. 1999, Ch. 4, p. 75, Gigerenzer & Goldstein). Gigerenzer et al. interpret these results as having radical implications for cognition in general – in particular, as undercutting the view that cognition must involve well-optimized cognitive machinery which behaves in accordance with classical rational norms of probability theory, logic, and decision theory. This line of thought raises the attractive possibility that the complexity of the mind may have been dramatically overestimated. Perhaps the mind is really just a collection of smart heuristics, rather than a fantastically powerful computing machine. This is an exciting and important thesis. This commentary focuses on three challenges to this approach, which may open up avenues for future research.

1. Empirical evidence. Gigerenzer et al. focus on providing a feasibility proof for the viability of a particular kind of simple reasoning heuristic. This task primarily involves providing computer simulations showing that simple heuristics give good results on specific decision problems, in comparison to conventional methods such as linear regression, and to other heuristic approaches, such as unit-weighted regression. But there is little by way of experimental evidence that people actually do reason in this way, aside from important but preliminary evidence reported in Chapter 7. This is particularly important precisely because the simulations in this book show that a wide range of algorithms give very similar levels of performance. Hence, prima facie, all these algorithms are equally plausible candidates as models of how people might perform on these problems.

In the absence of a broader set of experimental tests there is some reason to doubt that people make decisions by relying on one cue only. As Gigerenzer et al. note, in perception and language processing there is ample evidence that multiple cues are integrated in recognition and classification, in extremely complex ways (e.g., Massaro 1987). Gigerenzer et al. propose that these cases are in sharp contrast to the operation of conscious decision-making processes – determining whether this divide is a real one is an important area for empirical research.

2. Scope. One of the most startling findings in psychology is that, across a very wide range of judgment tasks, including medical diagnosis, expert performance does not exceed, and is frequently poorer than, results obtained by linear regression over sets of features of the cases under consideration (Meehl 1954; Sawyer 1966).

An equally startling finding, this time from artificial intelligence and cognitive science, has been that in everyday reasoning, people vastly outperform any existing computational model (Oaksford & Chater 1998a). Even the inferences involved in understanding a simple story draw on arbitrarily large amounts of world knowledge, and people must integrate and apply that knowledge highly effectively and rapidly. Attempts to model such processes computationally have become mired in the nest of difficulties known as the "frame problem" (Pylyshyn 1987).

So cognition is, in some regards, remarkably weak; and in other regards it is remarkably powerful. In the present context, the crucial point is that the simple heuristics discussed in this book are aimed at modeling areas where cognition is weak – indeed, where cognitive performance is already known to be frequently outperformed by linear regression. But it is by no means clear that the picture of the mind as a set of simple heuristics will generalize to everyday reasoning, where cognitive performance appears to be remarkably strong. Indeed, it may be that it is not that simple heuristics make us smart (as Gigerenzer et al.'s title suggests); rather it may be that we resort to simple heuristics to do the very thing we are *not* smart at.

3. Why *do heuristics work?* Gigerenzer et al. downplay the importance of traditional conceptions of rationality in their discussion of reasoning methods. Indeed, they note that a heuristic such as Take the Best has not been derived from "rational" principles of probability or statistics. Instead, they focus on an ecological notion of rationality – does the heuristic work in practice on real world data?

The viewpoint may appear to be an alternative to more traditional notions of rationality as used in psychology (Anderson 1990; Chater et al. 1999; Oaksford & Chater 1998b), economics (Kreps 1990) and behavioral ecology (McFarland & Houston 1981), in which behavior is assumed to approximate, to some degree, the dictates of rational theories, such as probability and decision theory. But it may be more appropriate to see the two viewpoints as complementary. Gigerenzer et al. (1999) are concerned to demonstrate rigorously *which* particular heuristics are successful, by computer simulation on realistic data sets. Traditional rational theories aim to explain *why* heuristics work. They characterize the optimization problem that the cognitive process, economic actor or animal faces; using rational theories (probability, decision theory, operations research) to determine the "rational" course of action; and conjecture that the heuristics used in actual performance approximate this rational standard to some degree. From this point of view, rational methods can be viewed as compatible with the "ecological" view of rationality outlined in Gigerenzer et al. (1999). Focusing on simple cognitive heuristics does not make the application of rational standards derived from formal calculi unnecessary. Instead, it gives a defined role for rational explanation – to explain why and under what conditions those heuristics succeed in the environment. This perspective is, indeed, exemplified in Gigerenzer et al.'s formal analysis of the conditions under which the Take the Best heuristic is effective (Ch. 6) and consistent with Gigerenzer et al.'s valuable comparisons between Take the Best and Bayesian algorithms (Ch. 8).

This book shows an important direction for research on human reasoning. It should act as a stimulus for empirical, computational, and theoretical developments in this area.

Simple heuristics could make us smart; but which heuristics do we apply when?

Richard Cooper

School of Psychology, Birkbeck College, University of London, London, WC1E 7HX United Kingdom **r.cooper@psychology.bbk.ac.uk**
www.psyc.bbk.ac.uk/staff/rc.html

Abstract: Simple heuristics are clearly powerful tools for making near optimal decisions, but evidence for their use in specific situations is weak. Gigerenzer et al. (1999) suggest a range of heuristics, but fail to address the question of which environmental or task cues might prompt the use of any specific heuristic. This failure compromises the falsifiability of the fast and frugal approach.

Gigerenzer, Todd, and the ABC Research Group (1999) are right to criticise much contemporary psychological decision-making research for its focus on mathematically optimal approaches whose application requires unbounded time and knowledge. They have clearly demonstrated that an agent can make effective decisions in a range of ecologically valid decision-making situations without recourse to omniscient or omnipotent demons. They have also cogently argued that biological decision-making agents cannot have recourse to such demons. The question is therefore not "Do such agents use heuristics?", but "Which heuristics do such agents use (and when do they use them)?" Gigerenzer et al. acknowledge that this question is important, but address it only in passing.

Gigerenzer et al.'s failure to specify conditions that might lead to the use of specific fast and frugal heuristics compromises the falsifiability of the fast and frugal approach. Difficult empirical results may be dismissed as resulting from the application of an as-yet-unidentified fast and frugal heuristic or the combination of items from the "adaptive toolbox" in a previously unidentified way. Gigerenzer et al. criticise the heuristics-and-biases approach of Tversky and Kahneman (1974) on much the same grounds. They note, for example, that both base-rate neglect and conservatism (two apparently opposing phenomena) can be "explained" by appealing to the appropriate heuristic or bias (because Tversky & Kahneman provide insufficient detail on the conditions that are held to evoke particular heuristics or biases). Gigerenzer et al. contend, quite reasonably, that such a post-hoc appeal does not amount to an adequate explanation of the behaviour. It is suggested that the use by the ABC Research Group of precise computational simulation techniques avoids this criticism. The computational simulations are very welcome and add a dimension often lacking in decision making research, but they do not, in their disembodied form, address the question of which heuristic might be applied when.

The issue of heuristic selection is not entirely ignored by Gigerenzer et al. The suggestion is that heuristics are either selected (or built from components within the adaptive toolbox) on a task-by-task basis. The challenge therefore lies in specifying: (1) the conditions under which different established heuristics are employed; (2) the conditions that provoke the construction of novel, task-specific heuristics; (3) the basic components available in the adaptive toolbox; and (4) the mechanisms by which appropriate heuristics may be constructed from these components. Of these, only the third is discussed at any length – the building blocks are held to involve elements that direct search, stop search, and decide based on the results of search – but that discussion of these elements is insufficient to allow the fourth concern to be addressed.

At several points in the book heuristic selection is conceived of as a meta-level decision-making task, suggesting that one might use a (presumably fast and frugal) heuristic to decide which fast and frugal heuristic to apply in a given situation. Two issues of minor concern are the possibility of an infinite regress (How do we select the fast and frugal heuristic to decide which fast and frugal heuristic to apply in the first place?) and the possibility that (assuming the infinite regress is avoided) the fast and frugal heuristic doing the selection might not yield the optimal fast and frugal heuristic for the original decision-making task.

Speculation as to the environmental and task cues that might lead to selection of the fast and frugal heuristics discussed in the book does not yield an obvious solution. For example, it is implied that one-reason decision making is appropriate (and presumably employed) for binary choice and that Categorisation By Elimination (CBE) is appropriate (and presumably employed) for multiple choice. However, both heuristics might be applied in both situations. The question of which heuristic to apply when remains.

The importance of heuristic selection is compounded by the lack of evidence presented in favour of the use by human decision makers of many of the heuristics discussed. For example, both QuickEst and CBE are presented solely as heuristics that can be shown to be fast and frugal. No comment is made on the psychological reality of either of these heuristics. This seems particularly odd when robust psychological findings that would appear to be of relevance (such as those addressed by Tversky and Kahneman's [1974] heuristics-and-biases approach) are ignored. How, for example, might the fast and frugal approach address the phenomena that Gigerenzer et al. use to demonstrate the difficulties present in the heuristics-and-biases approach (base-rate neglect and conservatism), or the confirmation bias often seen in diagnosis versions of categorisation? The latter, in particular, appears to be in direct conflict with the only categorisation heuristic proposed (CBE).

Fast and frugal heuristics have great promise. Human decision making cannot result from the application of algorithms with unbounded costs. Gigerenzer et al. have shown that fast and frugal heuristics can yield good decisions. They have not shown that humans use such heuristics, and by not addressing the question of which heuristics might be applied when, they have, like Tversky and Kahneman, given us a theory of human decision making that is unfalsifiable.