principles (*and,* for that matter, an array of rules of inference as well), is not, we daresay, very easy. Logicians can crack syllogisms in seconds, yes. But if you tried to configure your 403(b) in a thoroughly rigorous, decontextualized way, how long did it take you?

Other, arguably even deeper, problems spring from the simplicity of the problems that currently anchor the rationality debate. It seems bizarre to define general intelligence as the capacity to solve very easy problems. For example, Raven's Progressive Matrices, that vaunted "culture-free" gauge of *g*, can be mechanically solved (Carpenter et al. 1990). Once one assimilates and deploys the algorithm, does one suddenly become super-intelligent? Would a computer program able to run the algorithm and thereby instantly solve the problems, be counted genuinely intelligent? Hardly. (For more on this issue, see Bringsjord 2000. And recall Sternberg's continuous complaint that "being smart" in the ordinary sense has precious little to do with solving small, tightly defined test problems, a complaint communicated to some degree in his first-round commentary; cf. Sternberg 2000.)

Another problem arising from the fact that the rationality debate is tied to very easy problems is that psychology of reasoning is thereby structurally unable to articulate theories of robust human reasoning. Mental logic (championed, for example, by Rips 1994) cannot account for disproofs of the sort we gave above (because such disproofs are necessarily meta-proofs carried out outside a fixed set of inference schemas); and mental models theory (Johnson-Laird 1983), which rejects elaborate sequences of purely syntactic inferences, would seem to at least have a difficult time accounting for solutions to the problem we leave you with below (about which we've just given you a hint). What is needed is a theory of human reasoning that partakes of both the proof theoretic and semantic sides of symbolic logic, and the formal metatheory that bridges these two sides. (For a synoptic presentation of all this terrain, in connection to cognition and reasoning, see Bringsjord & Ferrucci 1998. For a theory of human reasoning designed to cover all of this terrain, Mental MetaLogic, see (Yang & Bringsjord, under review.)

Finally, what would be an example of a reasoning problem that *isn't* very easy, and the solving of which might justify confidence that the solver is both poised for success in the high-tech twenty-first century, and genuinely intelligent? Well, here's one; we refer to it as "The Bird Problem": Is the following statement true or false? Prove that you are correct.

(7) There exists something which is such that, if it's a bird, then everything is a bird.

Individual differences transcend the rationality debate

Elizabeth J. Newton[a] and Maxwell J. Roberts[b]

[a]*Department of Human Communication Science, University College London, London, WC1N 1PF, United Kingdom;* [b]*Department of Psychology, University of Essex, Colchester, Essex, CO4 3SQ, United Kingdom.*
liz.newton@ucl.ac.uk   mjr@essex.ac.uk

Abstract: Individual differences are indeed an important aid to our understanding of human cognition, but the importance of the rationality debate is open to question. An understanding of the process involved, and how and why differences occur, is fundamental to our understanding of human reasoning and decision making.

The main thesis of Stanovich & West (S&W) is that differences in individuals' performance can be used to cast light on the rationality debate. Even if we accept that this issue is important, and that humans occasionally behave irrationally, we still need clear criteria to identify such behaviour. Responses by themselves are often taken to be sufficient, but these are only informative if the cognitive processes underlying them are also understood. Otherwise,

there is little to gain by addressing the rationality question. This problem applies equally to the interpretation of psychometric test scores. Intelligence is a poorly understood construct, and the suggestion that it reflects only working memory capacity is by no means fully accepted.

Although highly intelligent people may be more likely to give normative responses in reasoning and decision-making tasks than less intelligent people, correlations between test score and reasoning performance can occur for a variety of reasons. Hence, differential correlations are not necessarily informative, and the focus on individual differences in terms of outputs rather than processes means that important qualitative differences are overlooked. The most straightforward reason for a correlation between intelligence test score and performance at a reasoning task is that highly intelligent people use the same processes as less intelligent people, but execute them more effectively. However, this merely leads us back to a "cognitive limitations" account of irrational behavior. Alternatively, perhaps highly intelligent people are better able to use *different,* more complex processes. A further possibility is that they are more likely to use *different* processes, but these are simpler, and hence more efficient. Either possibility leads into a debate about whether the strategy selections, rather than the responses themselves, are rational. A resolution depends crucially upon the ability to identify reasoning strategies accurately at the level of the individual. However, even where this is possible, we have to be certain that a suboptimal strategy, that is, one that is linked to poor performance, is really failing because of fundamental flaws. If, instead, a strategy is potentially normative, but too demanding to be executed accurately, then this turns the rationality issue into a debate concerning whether a person has made a strategy choice commensurate with his or her own ability to execute it accurately – and, ultimately, we are again returned to a cognitive limitations explanation of irrational behavior.

Given that people differ in the strategies they use, the rationality debate forces a dichotomy on us: are these choices normative or non-normative? Suppose you are presented with a series of trials, each consisting of compass point directions given together (e.g., one step north, one step east, one step north, one step west, one step south, one step south, one step west, one step north). The task is to determine the end point, relative to the start, after taking the steps. The natural strategy for this task is spatial: The full path is traced in the mind or by using a finger. For the cancellation strategy – a *task specific short-cut* – opposite steps are cancelled, with the remainder forming the correct response. Both strategies are normative: where applied accurately, they will yield the correct response. However, the spatial strategy is slower, less accurate and more demanding to execute. People are often painfully aware of the need to find an alternative. Surprisingly, even amongst university students, cancellation is used only by the minority. This is because it is only available to people with sufficiently high spatial ability to be able to identify the redundant processes of the spatial strategy and delete them, leading to the discovery of cancellation. Hence, people with high spatial ability outperform the rest, not because they are executing the spatial strategy more efficiently, nor because they are better able to use an enhanced spatial strategy in which additional processes increase accuracy, but because they have dispensed with spatial representations altogether, increasing accuracy and minimising effort (see Newton & Roberts 2000; Roberts et al. 1997).

So, are the spatial strategy users irrational for this task? Granted, if we focus on outputs only, they are less accurate than cancellation users; but, as suggested earlier, the rationality debate is only served crudely in this way. Errors when using the spatial strategy are due to capacity limitations in any case. The rationality debate is not served at all by considering whether selected strategies are normative – both will yield correct responses if executed accurately. Are the spatial strategy users less rational because they made an inappropriate choice? No, there is no alternative available to them, and hence no choice. If the

spatial strategy users must be branded as irrational, this can only be because they lack the necessary resources to discover the more efficient method. This conceptualisation of irrationality as a lack of creativity may be unappealing to some, but for this example the categorisation of these people as irrational is artificial and driven only by the perceived need to address the rationality question.

Overall, S&W make many valid points. We agree that individual differences are an important aspect of human cognition. But to use them merely to resolve the rationality debate is problematic and neglects their full potential. Issues of how people reason, and how these processes change and develop with experience can be better answered by not being side-tracked in this way.

Authors' Response

The rationality debate as a progressive research program

Keith E. Stanovich[a] and Richard F. West[b]

[a]*Department of Human Development and Applied Psychology, University of Toronto, Toronto, Ontario, M5S 1V6, Canada;* [b]*School of Psychology, James Madison University, Harrisonburg, VA 22807.* **kstanovich@oise.utoronto.ca   westrf@imu.edu   falcom.jum.edu//~westrf**

Abstract: We did not, as Brakel & Shevrin imply, intend to classify either System 1 or System 2 as rational or irrational. Instrumental rationality is assessed at the organismic level, not at the subpersonal level. Thus, neither System 1 nor System 2 are themselves inherently rational or irrational. Also, that genetic fitness and instrumental rationality are not to be equated was a major theme in our target article. We disagree with Bringsjord & Yang's point that the tasks used in the heuristics and biases literature are easy. Bringsjord & Yang too readily conflate the *ability* to utilize a principle of rational choice with the *disposition* to do so. Thus, they undervalue tasks in the cognitive science literature that compellingly reveal difficulties with the latter. We agree with Newton & Roberts that models at the algorithmic level of analysis are crucial, but we disagree with their implication that attention to issues of rationality at the intentional level of analysis impedes work at the algorithmic level of analysis.

We found much that is congenial to our way of thinking in these three commentaries. For example, we welcome **Brakel & Shevrin**'s points about Freud and dual-process theorizing. Such theorizing of course predates Freud as well, going back at least to Plato. As Plato writes in *The Republic,*

> we may call that part of the soul whereby it reflects, rational; and the other, with which it feels hunger and thirst and is distracted by sexual passion and all the other desires, we will call irrational appetite, associated with pleasure in the replenishment of certain wants. (Cornford 1945, p. 137)

While we welcome Brakel & Shevrin's addendum, we take it as understood that the purpose of our paper was not the historical exegesis of dual-process notions. Some historian really does need to do a treatise tracing dual process ideas from Plato, through Freud, to the cognitive revolution (e.g., Evans & Wason 1976; Shiffrin & Schneider 1977), but this was not our purpose. Our argument depended only upon common assumptions of these theories and not on nuanced differences or historical relationships.

There is much in **Brakel & Shevrin**'s characterization of System 1 and System 2 that we agree with. For example, we agree (as do most of the dual-process theorists that we cite in the target article) that System 1 processing is not supplanted by System 2 processing with development, but rather, that both types of processing continue to operate in parallel. However, there are some points of misinterpretation as well. Brakel & Shevrin seem to imply that we are labelling *systems* as rational or irrational, but this is not the case. Instrumental rationality (what was termed normative rationality in our target article) is assessed at the organismic level, not at the subpersonal level. Neither System 1 nor System 2 are themselves inherently rational or irrational.

We focused in the target article and elsewhere (e.g., Stanovich & West 2003) on situations where System 1 functioning served to disrupt the pursuit of instrumental rationality (if not overridden by System 2 processes). But we were also clear to note in the target article that "It must be stressed though that in the vast majority of mundane situations, the evolutionary rationality embodied in System 1 processes will *also* serve the goals of normative rationality" (Stanovich & West 2000, p. 661); and in our Authors' Response we repeated that "we made it clear in the target article that in most cases the goals of Systems 1 and 2 will coincide and that System 1 processes will often *also* serve the goal of normative rationality" (p. 708). So, System 1 serves the organism most of the time by facilitating instrumental rationality, but sometimes disrupts the pursuit of instrumental rationality and must be overridden by System 2.

Thus, System 1 is not appropriately characterized *itself* as being either inherently rational or irrational – a point we feel we made clear in the original target article. Furthermore, the same is true of System 2. It can instantiate rules of rational thought which facilitate maximal goal satisfaction (our emphasis in the target article), but it can also instantiate ideas and rules (memes, in the view of Dennett 1991 and Blackmore 1999) that impede the organism's pursuit of instrumental rationality – a theme we did not emphasize in the target article, but have stressed in subsequent publications (Stanovich 2004; Stanovich & West 2003). Thus, System 2 likewise should not be characterized as inherently rational or irrational, since it too is a subpersonal entity.

Brakel & Shevrin seem to have been confused by our use of the term evolutionary rationality, but here the fault might be ours. Our use of the term in the target article was perhaps too clever by half. The term was coined as an indirect tweak at the evolutionary psychologists who conflate behavior serving genetic fitness with behavior that is instrumentally rational (a major theme in our book-length treatments of these issues; Stanovich 1999; 2004). The terms evolutionary rationality (behavior serving genetic fitness) and normative rationality (instrumental rationality) were meant to separate these two. For example, Over (2000), in his critique of work on fast and frugal heuristics (e.g., Todd & Gigerenzer 2000), makes use of our distinction in exactly the way we intended. Nevertheless, we acknowledge that the term evolutionary rationality may have invited people to conflate just the distinction that we wished to emphasize (as Brakel & Shevrin seem to have done). Thus, in a new book by one of us (Stanovich 2004) – which is largely devoted to working out the implications of mismatches between behavior serving the interests of replicators in the environment of evolutionary adaptation and current instrumental rationality for the organism – the term