important aspect concerns the developmental dimension. The primary process is called primary because it precedes the secondary process in development. According to some of our own evidence, the primary process appears to prevail until about seven years of age (Brakel et al. 2002). However, even though primary process is then largely replaced by secondary process, the primary process continues to operate unconsciously and may appear consciously under circumstances of stress and conflict. (Brakel & Shevrin, submitted a; submitted b) From a Freudian standpoint, the primary and secondary process systems are co-existent throughout adult life.

The second matter concerns whether the primary process, and the other System 1 operations, need to be considered "rational" in any way – including "evolutionarily rational." We would like to suggest that, (1) System 1 operations, like the primary process, are neither rational nor irrational, but instead *a*-rational; (2) nonetheless, both the a-rational System 1 and the rational System 2 can each have different roles in enhancing evolutionary fitness; (3) System 1 (primary process type) operations are deemed "rational" (incorrectly, in our view) precisely when their operations effect increased evolutionary fitness.

(1) The a-rationality of System 1 can be seen in a certain category mistake made by birds, frogs, and people. A bird responds to a big black cloth just as it does to a bird bigger than itself. A frog swallows BB pellets with the same alacrity as it does bugs and flies. A spider-phobic person responds with dread to a plastic or rubber replica of a spider much as he/she does to a real spider. In all of these instances we are not dealing merely with ambiguous or unclear perceptions. In each instance, the black cloth, BB pellet, and spider replica, even when seen up close and (certainly in the human case and likely in the others) accurately perceived, still occasion the fearful flight, swallowing, and dread reactions. We are dealing with "funny System 1 categories" – something to flee from, something to swallow, something to fear – that include items resembling the truly biologically relevant exemplars only superficially. Individual features, often nonessential ones, determine category membership. Now, while it could be argued that in the human case the individual is aware of the seeming irrationality of his/her response, nevertheless the fearful response to the replica spider persists based on the a-rational category match. Unlike people, birds are not rational and therefore cannot be irrational. Birds make no rational appraisals of the a-rational category matches underlying their urgent flights from big black cloths. Where there is no capacity to choose in a rational fashion, behavior predicated on a-rational principles is not irrational.

(2) Both System 1 and System 2 can play different roles in evolutionary fitness. There is general agreement that employing the basic rules of logic, especially in the service of reality testing, can be useful to individuals who are capable of such rational cognitive feats. Indeed, as pointed out in the target article, such capacities certainly can further individual goals. But clearly, insofar as System 2 operations are invaluable, not just in the psychology lab but in obtaining food, shelter, and mates, having these capacities will also enhance the reproductive fitness of individuals. With respect to System 1 operations, these too can enhance evolutionary success – but precisely because System 1 is not a rational system and therefore does not have the constraints of rationality.

Suppose bugs, whose life cycles are much faster than those of frogs, evolve forms that weigh as much as metal BBs and even taste like them. Frogs behaving on the basis of their System 1 category will eat better and reproduce more successfully than would frogs who could make System 2 inferences like, "If round and weighty and metallic tasting, then not food, don't swallow." Likewise, the bird with the "funny category" titled "flee, it is big!" will be able to avoid more novel non-avian-but-nonetheless-dangerous-things. Further, unless useless flights are too much of an energy drain, the birds with a-rational categories will do better reproductively than "rational" birds, who would, for example, fail to avoid dangerous human-made items such as pieces of steel hoisted skyward in a construction area. Even if useless flights compromise reproductive potential, they are still less damaging than sudden death. In the human case, to the extent that spiders are truly dangerous, it is best to avoid anything that might possibly be a spider – even if the judgment is made under the quick and dirty System 1 – since spider bites can be fatal, and avoiding rubber spiders merely silly.

The phenomenon of stranger anxiety provides a more general example of a possible adaptational advantage of a-rational mentation in people. Although stranger anxiety does not occur in all healthy babies, it does occur in many at around eight months. Infants younger than eight months, largely System 1 thinkers, can be comfortably cared for by any nurturing adult. Not so when certain babies become more capable of even rudimentary System 2 categorizing and infer, "since these are not my parents, I should be wary." (For further examination of the evolutionary role of a-rational mentation, see Brakel 2002).

Note here that both of our examples concerning people demonstrate that evolutionary explanations can provide only necessary, but not sufficient conditions for a-rational mediated behaviors in human beings. On the basis of temperament, some normal babies may have a very mellow version of wariness. And with regard to spider phobias, indeed spiders, rather than butterflies, may have been selected by evolution as a target for phobias; but for each spider phobic there must also be specific psychological experiences that result in the development of a phobia.

(3) When a-rational mechanisms yield success, it is not on the basis of rationality, but because they function with fewer constraints and more category inclusiveness. While we are not claiming that these System 1 operations work most of the time – many pointless flights, nutrition-less swallowings take place – they work just often enough to enhance reproductive fitness and evolutionary success. But when this outcome is thought to demonstrate "evolutionary rationality," more than just a misnomer is involved. It is precisely the "evolutionary a-rationality" that has contributed to enhanced reproductive fitness. In other words, System 1 operations, whether or not successful, do not result from irrationality or error in an otherwise rational system. Rather, the evolutionary utility of System 1 is due to precisely its nonrational or arational principles – principles operating quite differently from System 2 rationality.

As we stated above, Freud posited that although there is developmental progression whereby primary process is largely supplanted by secondary process, primary process continues to operate, often unconsciously, often influencing behavior significantly. Indeed, for Freud, primary process and secondary process systems are co-existent throughout adult life. Now more than a century later, as it becomes clearer that a-rational System 1 and rational System 2 are each important for enhancing evolutionary fitness in different ways, Freud's early conceptualizations of dual processing systems look better and better.

The problems that generate the rationality debate are too easy, given what our economy now demands

Selmer Bringsjord and Yingrui Yang

Department of Philosophy, The Minds and Machines Laboratory, Psychology and Cognitive Science, and Department of Computer Science, Rensselaer Polytechnic Institute (RPI), Troy, NY 12180; **{selmer;yangyri}@rpi.edu**
www.rpi.edu/~brings

Abstract: Stanovich & West (S&W), following all relevant others, define the rationality debate in terms of human performance on certain well-known problems. Unfortunately, these problems are very easy. For that reason, if System 2 cognition is identified with the capacity to solve them, such cognition will not enable humans to meet the cognitive demands of our technological society. Other profound issues arise as well.

The rationality debate revolves around a set of problems, nearly all of which, of course, are well known to the participants in this

https://doi.org/10.1017/S0140525X03220112 Published online by Cambridge University Press

debate. But all these problems are, to put it bluntly, very easy. This fact – to which the researchers who have hitherto defined the debate are apparently oblivious – has far-reaching consequences, as we begin to explain in this commentary.

To save space, we focus here upon deductive reasoning, and specifically upon syllogistic reasoning. We label a logic problem as "very easy" if there is a simple, easily taught algorithm which, when followed, guarantees a solution to the problem. Normal cognizers who take an appropriate first course in symbolic logic can master this algorithm: Represent a syllogism in accordance with Aristotle's A/E/I/O sentences, cast this representation in first-order logic (FOL), inspect the formalization to see if a proof is possible, carry out the proof if it is, or carry out, in accordance with a certain sub-algorithm, a *dis*proof if it isn't. For 14 years, year in and year out, Bringsjord's students have achieved a more than 95% success rate on post-tests given in his "Introduction to Symbolic Logic" course, in which they are asked to determine whether or not syllogisms are valid. This includes syllogisms of the sort that S&W report subjects to be befuddled by. As an example, consider the "challenging" syllogism S&W present:

(1) All mammals walk.
(2) Whales are mammals.
Therefore: (3) Whales walk.

Each of these sentences is an A-sentence (All A are B):

(1′) All M are A.
(2′) All W are M.
Therefore: (3′) All W are A.

So in FOL we have:

(1″) $\forall x\,(Mx \rightarrow Ax)$
(read: for all x, if x is an M, then x is an A)
(2″) $\forall x\,(Wx \rightarrow Mx)$
Therefore: (3″) $\forall x\,(Wx \rightarrow Ax)$

The proof now runs as follows: Let *a* be an arbitrary thing. We can instantiate the quantifiers in (1″) and (2″) to infer $Ma \rightarrow Aa$ and W*a* (M*a*), respectively. We can then use hypothetical syllogism (a "chain rule") to conclude $Wa \rightarrow Aa$. Since *a* was arbitrary, from this we can conclude by universal introduction $\forall x\,(Wx \rightarrow Ax)$. QED.

For every formally valid syllogism, the corresponding proof can be generated by such simple mechanical means. What about formally invalid syllogisms? Producing disproofs is here once again a matter of following a trivial algorithm. To show this, consider an example from Johnson-Laird & Savary (1995). When asked what can be (correctly) inferred from the two propositions

(4) All the Frenchmen in the room are wine-drinkers.
(5) Some of the wine-drinkers in the room are gourmets.

most subjects respond with

Therefore: (6) Some of the Frenchmen in the room are gourmets.

Alas, (6) cannot be derived from (4) and (5), as can be seen by inspection after the problem is decontextualized into FOL, and chaining is sought.

But Bringsjord's students, trained to use both the algorithm above, and therefore the sub-algorithm within it for generating disproofs, and nothing else, not only cannot make the erroneous inference, but *can also* prove that the inference is erroneous. Here's why. The Aristotelean form consists of one A-sentence and two E- sentences (Some A are B):

(4′) All F are W.
(5′) Some W are G.
Therefore: (6′) Some F are G.

In FOL this becomes

(4″) $\forall x\,(Fx \rightarrow Wx)$
(5″) $\exists x\,(Wx\ \&\ Gx)$
Therefore: (6″) $\exists x\,(Fx\ \&\ Gx)$

Notice, first, that neither Wa nor Ga can be used to chain through $Fa \rightarrow Wa$ to obtain the needed Fa. Next, for a disproof, imagine worlds whose only inhabitants can be simple geometric shapes of three kinds: dodecahedrons (dodecs), cubes, and tetrahedrons (tets). Suppose now that we fix a world populated by two happy, small dodecs, two happy, large cubes, and two medium tets. In this world, all dodecs are happy (satisfying premise [4″]), there exists at least one happy, large thing (satisfying premise [5″]), and yet it is not the case that there is a large dodec (falsifying proposition [6″]). Students in Bringsjord's logic course, and in logic courses across the world, mechanically produce these disproofs, often by using two software systems that allow for such worlds to be systematically created with point-and-click ease. (The systems are Hyperproof and Tarski's World, both due to Barwise & Etchemendy 1984; 1999.) One of us has elsewhere argued that the appropriate pedagogical deployment of these two remarkable systems substantiates in no small part the neo-Piagetian claim that normal, suitably educated cognizers are masters of more than System 2 cognition at the level of FOL (Bringsjord et al. 1998). Whether or not Bringsjord is right, it's hard to see how S&W consider the neo-Piagetian response to the normative/descriptive gap. They consider a quartet of proposed explanations – fundamental irrationality, performance errors, computational limitations, misconstrual of problem. But why can't the gap be explained by the fact that most people are just uneducated? (In his first-round commentary, Zizzo [2000] mentions the possibility of teaching logic on a mass scale, but then seems to reject the idea. Actually, by our lights, that's exactly what needs to be done in order to meet the demands of our high-tech economy.)

Now we know that S&W, in responding to Schneider's (2000) first-round commentary, point out that the correlation between heuristics and biases tasks and training in mathematics and statistics is negligible (Stanovich & West 2000, p. 705). But this is irrelevant, for two reasons. First, S&W ignore Schneider's specific claim about syllogisms, and (tendentiously?) zero in on her claim that suitable education can cultivate a cognition that leads to higher SAT scores. What Schneider says about syllogisms is that some people can effortlessly and accurately assess them (albeit via System 1 cognition in her cited cases). Second, the issue, in general, is whether *specific* training has an effect on performance. Few math courses (traditionally, none before analysis) at the undergraduate (and even, in more applied departments, at the graduate) level explicitly teach formal deductive reasoning, and many first logic courses are merely courses in *informal* reasoning and so-called critical thinking – courses, therefore, that don't aim to teach decontextualization into some logical system. This is probably why the problem of moving from mere problem solving in mathematics to formal deductive reasoning (a problem known as "transition to proof"; Moore 1994) plagues nearly all students of math, however high their standardized test scores; and why, in general, there is little correlation between math education and the solving of those problems in the rationality debate calling for deductive reasoning. The meaningful correlation would be between subjects who have had two or more courses in symbolic logic and high performance, for example, on (very easy) deductive reasoning problems seen in the rationality debate. We predict that this correlation will be strikingly high. (See also the prediction made by Jou [2000, p. 680] in the first round of commentary, concerning scores on the logical reasoning section of the GRE and normative performance. In this connection, it is probably noteworthy that those who write on logical reasoning in "high stakes" standardized tests invariably have training in symbolic logic.)

We heartily agree with S&W that today's workforce demands rigorous, deoncontextualized thinking on the part of those who would prosper in it. In their response to the first round of commentaries, the authors provide a nice list of relevant challenges (p. 714); let's take just one: deciding how to apportion retirement savings. In our cases, which are doubtless representative, we can choose to set up our 403(b)'s with one of three companies, each of which offers, on the mutual fund front alone, one hundred or so options. One tiny decision made by one fund manager makes syllogistic reasoning look ridiculously simple by comparison, as any of the proofs driving financial knowledge-based expert systems make plain. To assess the future performance of many such managers making thousands of decisions on the basis of tens of thousands of data points, and at least hundreds of declarative 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

https://doi.org/10.1017/S0140525X03220112 Published online by Cambridge University Press