Rationality Disputes – Psychology and Epistemology

Patrick Rysiew*
University of Victoria

Abstract

This paper reviews the largely psychological literature surrounding apparent failures of human rationality (sometimes referred to as 'the Rationality Wars') and locates it with respect to concepts and issues within more traditional epistemological inquiry. The goal is to bridge the gap between these two large and typically disconnected literatures – concerning rationality and the psychology of human reasoning, on the one hand, and epistemological theories of justified or rational belief, on the other – and to do so in such as way as to expose interesting points of contact and convergence between them.

1. Introduction

In addition to being of central interest within traditional epistemology, and among our terms of highest intellectual praise, *rationality* is central to our self-conception: Aristotle held that we are '*rational animals*', a presumption built into the very name of our species ('*homo sapiens*'); and the thought that humans are rational, perhaps distinctively so, appears to be part of the popular fabric of thought about ourselves. There is long-standing disagreement among epistemologists as to the nature of epistemic rationality (hereinafter, simply 'rationality') – this is a matter, in fact, that is as controversial as the nature of justification and knowledge. There has also recently arisen heated debate – sometimes referred to as 'the Rationality Wars' – among psychologists and philosophers of psychology concerning what we should say in the face of empirical findings about humans' apparently disappointing performance on certain 'reasoning tasks'. According to some, those results force us to confront the possibility that humans may not be so rational after all – that they may in fact be quite irrational; according to others, such results, together with a psychologically realistic view of how human reasoning[1] actually proceeds, point up the need to revise standard views of what rationality involves. This paper briefly reviews 'the Rationality Wars' and identifies points of contact between that debate and concepts and issues within more traditional epistemological inquiry. The goal is to bridge the gap

between these two large and typically disconnected literatures – concerning rationality and the psychology of human reasoning, on the one hand, and epistemological theories of knowledge and justified/rational belief, on the other – and to do so in such as way as to expose interesting points of contact and commonality between them.

2. Some Preliminaries: Epistemology and Rational Belief

Following Aristotle's classification of the virtues, it is common to distinguish between 'epistemic' and 'practical' rationality. In either case, there is an important normative question to be addressed. One common way of marking the distinction is in terms of whether it is beliefs or actions which are being assessed. But this won't do, since beliefs can be assessed in terms of their promoting the achievement of one's patently non-epistemic goals (see, e.g., Bonjour, *Structure of Empirical Knowledge* 6; Harman, 'Rationality' §1.2; 'Practical Aspects'). This, however, points to a better marker of the distinction we're after: epistemic rationality (and epistemic reasons, etc.) has to do with the effective pursuit of the distinctively cognitive-epistemic end of true belief (Bonjour, *Structure of Empirical Knowledge* 6; Harman, 'Rationality' §1.2; 'Practical Aspects'; see too, e.g., Audi 17–18; Goldman, 'Unity of the Epistemic Virtues'; Nozick 71). Not that there are no questions confronting such a *veritistic* orientation: for instance, one might wonder whether the goal of acquiring true beliefs doesn't require too the (distinct?) goal of avoiding false ones, not to mention the further (sub-?) goals of acquiring beliefs promoting the satisfaction of familiar criteria of theory choice, such as simplicity, explanatory power, and so on; more dramatically – since *it* is obviously an *alternative* to a thoroughly veritistic approach – one might hold, e.g., that there are some non-truth-linked epistemic ends and/or goods.

Such questions aside, however, it is fair to say that the majority view within epistemology is that whether it is beliefs, their possessors, or the processes by which beliefs are formed that we are assessing, it is a connection with truth which marks off the distinctive nature of epistemic rationality. Further, within epistemology, particularly among those of an 'internalist' stripe (on which, see below), 'rational', as applied to beliefs, is often taken to be synonymous with 'justified' (e.g., S. Cohen, 'Justification and Truth' 283; Cruz and Pollock 1215 n. 1) – in fact (again, particularly among internalists) it is commonly taken to designate the very sort of justification that is supposed to convert ungettiered[2] true belief into knowledge.

Of course, as we'll see below, the proper way of conceiving of the connection between justification (rationality) and truth is the subject of great controversy. Nonetheless, such disputes typically take place against the background of the following general view: rationality, or rational processes of belief-acquisition, are in one way or another (see below) proper means of pursuing the distinctively cognitive-epistemic goal of

truth; rational beliefs are those formed or backed by such processes; and rational agents are those whose beliefs are typically thus formed. Approaching things in this way, rationality emerges as *the* central epistemological notion: beliefs and belief-forming processes *per se* are the province of psychology, and truth is the province of metaphysics (or semantics, according to some), leaving rational/justified belief as the sole distinctively epistemic component in knowledge and the primary marker of epistemically laudable belief (see Kim 382–3; cf. Bonjour, *Structure of Empirical Knowledge* 5).

Just as much a part of mainstream epistemology as the centrality of justification or rationality is a broadly anti-sceptical orientation – i.e., a belief that we know many things, (therefore) that many of our beliefs are justified or rational, and that we, considered as epistemic agents, are (therefore) for the most part rational.[3] (And perhaps distinctively so, since, it is thought, many non-human animals may lack certain concepts or abilities required for rational belief.) Not that anyone denies the existence of occasional lapses of rationality – an inattentiveness to the evidence, indulging in wishful thinking, and so on. But to the extent that one shares the general anti-sceptical orientation of most epistemologists, such happenings will be regarded as exceptions occurring on the periphery, so to speak, of our cognitive-epistemic lives. That we *are* – for the most part, and for much of the time – rational animals, is not a mere Aristotelian conceit; nor is it – like, say, the belief that the comparative intelligence and success of humans is directly attributable to larger brain size – merely a piece of easily dispelled 'popular mythology'; it is a perhaps indispensable starting point for epistemological theorizing.[4] And yet, in recent decades, a stream of experimental results has threatened to show that *homo sapiens* may not be so sapient after all.

3. Irrational Animals? Some Experimental Results[5]

If we are indeed rational, and so capable of adopting proper and effective means for pursuing the goal of truth, it is very natural to suppose that we must have a grasp of certain formal rules governing logical and probabilistic inference. Such a grasp, in fact, is sometimes thought to be at the very core of rationality, ensuring as it does (the thinking goes) the proper connection between one's evidence and beliefs, as well as the appropriate sort of consistency amongst the latter. Thoughts such as these lie behind a particular view of rationality, which has been called the 'Classical' conception (Chase, Hertwig, and Gigerenzer) – or, because the apparently widespread acceptance it enjoys, simply 'the Standard Picture' (SP):

> According to this picture, to be rational is to reason in accordance with principles of reasoning that are based on rules of logic, probability theory and

so forth. If the standard picture of reasoning [rationality] is right, principles of reasoning that are based on such rules are normative principles of reasoning, namely they are the principles we ought to reason in accordance with. (Stein 4)

In spite of the broad popularity it enjoys, however – but also in part because of it! –, SP has been subjected to serious scrutiny of late. For, when coupled with various well-known experimental findings, the Classical conception or SP appears to have the consequence that we are not rational after all.

Perhaps the best-known such experiment is the Wason Selection Task ('Reasoning about a Rule'). In this study, designed to test people's facility with rules of deductive inference, subjects are given the following task:

Here are four cards. Each of them has a letter on one side and a number on the other side. Two are shown with the letter side up, two with the number side up

| E | C | 5 | 4 |

Which of the cards do you have to turn to determine whether the following claim is true: *If a card has a vowel on one side, then it has an odd number on the other side*?

The right answer here, of course, is that the 'E' and '4' cards must be turned over, since only they constitute possible counter-instances to the claim in question. However, most subjects answer that the 'E and '5' cards need to be turned over. Once again: the '5' response is incorrect because, regardless of what is on the other side, the card does not represent a possible counter-instance to the claim in question. (The latter tells us, 'if vowel, then odd number', not, 'no odd number unless vowel'.) Thus, subjects seem to exhibit 'confirmation bias'; formally speaking, they seem either not to grasp, or to not apply correctly, *modus tollens*.

Bad as that may seem, there is equally disquieting news concerning ordinary subjects' performance when it comes to probabilistic reasoning. Thus, Tversky and Kahneman ('Judgments of and by Representativeness' 91–3) presented subjects with a description of some fictional person:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issue of discrimination and social justice, and also participated in anti-nuclear demonstrations.

They then asked subjects to rank the following statements from most to least probable:

(a) Linda is a teacher in elementary school
(b) Linda works in a bookstore and takes yoga classes

(c) Linda is active in the feminist movement
(d) Linda is a psychiatric social worker
(e) Linda is a member of the League of Women Voters
(f) Linda is a bank teller
(g) Linda is an insurance sales person
(h) Linda is a bank teller and is active in the feminist movement.

Eighty-nine percent of 'naïve' subjects (undergraduates without any background in probability and statistics) rated (h) more probable than (f). Among 'sophisticated' subjects (graduate students in the Decision Science Program of the Stanford Business School) 85% still gave the same answer. But it's a very basic rule of the probability calculus that a conjunction (*a* & *b*) cannot be more probable than either of its conjuncts. Thus: Linda can't be a *feminist* bank teller unless she's a *bank teller*; (f), therefore, is at least as probable as (h). But even non-naïve subjects judge to the contrary, thereby committing *the conjunction fallacy*.

These are just two of the better-known examples of studies seeming to show that people fail to reason as SP says they should. But they're by no means the only such results. Thus, for instance, people regularly commit *the gambler's fallacy*, supposing that the probability of some event (e.g., that a fair coin will come up heads next time it's tossed) is somehow influenced by a certain pattern of previous occurrences (its having come up tails the previous five tosses). So too, there is *base-rate neglect*, a clear violation of Bayesian probability calculus: even those whom one would hope to have been trained in such matters – e.g., medical staff and students, asked to estimate the probability that a given patient has some disease, given a positive result on a test with such-a-such incidence of false positives – ignore information they're explicitly presented with about the prevalence of the disease in the general population (Casscalls, Schoenberger, and Graboys).

More examples still could be given – of, e.g., *belief perseverance* (Nisbett and Ross), *overconfidence bias* (see Lichtenstein, Fischoff, and Phillips), *belief bias* (see Evans, 'In Two Minds'), *anchoring* (Tversky and Kahneman, 'Judgment under Uncertainty'), and so on. But enough has already been said to make it at least understandable why, according to some – for instance, the psychologists Nisbett and Borgida – the prospects for human rationality might be rather 'bleak' (935). Humans, as such theorists see it, are simply deeply irrational. Rather than right rules of logic, probability theory, and so on, people reason – or, at least, they *think* – using a grab-bag of various '*heuristics*' of the sort that Kahneman et al. have described, heuristics which bring in their train certain illogical '*biases*'. For instance, where subjects commit the conjunction fallacy, it is plausible to see them as employing 'the representativeness heuristic' – (roughly,) *to judge the probability that an object, X, belongs in some category, C, look at the degree to which x is similar to, or representative of, typical members*

of C. Hence, for instance, the erroneous verdict that it is more probable that Linda is a feminist bank teller than that she is a bank teller.

4. *Some Philosophical Responses:* Must *Humans be Rational?*

The discussion of the previous two sections nicely recapitulates what Botterill and Carruthers (105) have described as a near-disciplinary divide that has until recently prevailed, with philosophers confident of human rationality, and psychologists finding only grounds for pessimism about it. More recently, however, various psychologists have staked out a more positive take on the relevant data, one which promises to preserve the picture of humans as rational. But before turning to such views, as we will in the next section, it's worth pointing out one respect in which the rather pessimistic assessment of human reasoning just described is surely too gloomy.[6] For, after all, even if we regularly fail to be SP-rational, we – or some of us, anyway – can come to recognize that fact and mend our ways, so that we're not hostage to those fallacious inferences to which we're naturally prone. So we seem, as a matter of fact, to be at least *capable* of reasoning in a manner which accords with SP. This, as much as our committing the errors that we do, is an important datum that needs explaining. Still, several prominent philosophical responses to the discussion of the last section purport to show something rather stronger: namely, that, for one reason or another, we humans *can't* be irrational, or that we literally cannot make sense of the suggestion that we are.

First, there is the *interpretationist* argument, whereby regarding others as (largely) rational is a precondition of crediting them with intentional states and engaging in reasoning at all (e.g., Davidson; Dennett). Second, there is the argument from *reflective equilibrium* (e.g., L.J. Cohen), according to which our views concerning rationality should be arrived at through a process of mutual adjustment between our intuitions about the correctness (or not) of particular inferences and our assessment of proposed general principles of rationality, with the aim being to arrive at an acceptable systematization of the latter. In this back and forth, there is no higher court of appeal than humans' intuitions about various cases and principles; as Cohen puts it, 'a normative theory [of rationality] must be based ultimately on the data of human intuition' (321). The result is that such reasoning mistakes as we commit *must* be regarded as mere performance errors, which do not impugn the reasoning *competence* which all normal humans possess (321). Third, there is the *evolutionary* argument, which states that we must be (largely) rational, since our ancestors, had they been *il*logical, would not have survived to reproduce.

Each of these arguments has received substantial critical attention elsewhere,[7] and the consensus seems to be that none obviously succeeds in establishing significant in-principle limitations on the extent to which

humans might be found to be irrational. But a brief discussion of these arguments sets the stage for a different type of response to any very pessimistic assessment of human rationality, one which calls into question the aptness of the normative (SP) standard being used is arriving at such assessments.

First, then, as to the *interpretationist* argument, even if intentional ascription requires that we regard the subject as respecting *some* of the inference patterns singled out by SP (see Cherniak), there is no reason to suppose that it will require that we regard the subject as respecting all, or even most, of them. After all, charitably interpreted, the principle of charity (Quine, *Word and Object*), say, enjoins us to avoid attributing *inexplicable* error to others; but, as we saw just above, the use of various illogical 'rules of thumb' of the sort posited by those in the 'heuristics and biases' tradition makes our failure to conform to SP all too easy to understand.

Similar points can be made concerning the argument from *reflective equilibrium*. Since the gambler's fallacy, for instance, is likely to be in reflective equilibrium for most ordinary subjects, if we want to foreclose the possibility that the principle underlying it will come out as normatively correct, the sort of reflective equilibrium that should be taken as relevant to theorizing about rationality is *expert wide reflective equilibrium*, which would have among its inputs the results of both scientific studies of our cognitive capacities and natural belief-forming tendencies, and disciplines such as statistics, etc. (Stein ch. 5; Stich ch. 4; Stich and Nisbett). And while it's generally agreed that facts about the finiteness of human cognitive resources, for example, should inform what can reasonably be required for rational belief, there is no reason why such wide expert reflective equilibrium would *have* to issue in an assessment of ourselves as extremely, much less perfectly, rational. So too, the fact that we – or some of us – can recognize various reasoning errors *as* errors (see above) needn't be taken to show that humans generally must already possess 'tacit knowledge' of the relevant (correct) rules; perhaps – to preview an idea to which we'll return below – it is merely that, thanks to their having had the appropriate education and training, say, some have been able to acquire such knowledge.

Lastly, there is the *evolutionary* argument, expressions of which abound. For instance:

> Natural selection guarantees that organisms either know the elements of logic or become posthumous. (Fodor 121)
> Natural selection guarantees that most of an organism's beliefs will be true, most of its strategies rational. (Dennett 75)
> It would be a cruel hoax and evolutionarily disadvantageous for nature to provide us with propositional representations if our ways of connecting them and reasoning with them failed to preserve their propositional status. Because propositions are profoundly truth functional, their inference procedures ought

to be truth preserving . . . Logicians refer to this property as logical soundness. (O'Brien 113)[8]

Here, the idea is that a terribly pessimistic construal of the experimental evidence *can't* be right, since natural selection would have shaped our ancestors' – hence, *our* – minds in such a way as to assure their – and our – rationality. In effect, we wouldn't be here wondering whether we're rational *unless we were*!

Closely related to the evolutionary argument against the possibility of widespread human irrationality are statements of the so-called 'rationality paradox' (Evans and Over 2–4): in many ways, humans appear to be a remarkably successful and intelligent species; and yet, *just look at those experimental results!* In the words of one psychologist, 'We can put a man on the moon, so why can't we solve those logical reasoning problems?' (O'Brien 110).

Why not regard the results of the reasoning experiments simply as interesting empirical discoveries? Why think they constitute one half of a *paradox*? Not *merely* because we had thought that we were rational, but because here, as in the evolutionary argument, it is being assumed that what's responsible for our species' success and accomplishments *are precisely the sorts of SP-type skills that the reasoning experiments test*. But this assumption can be challenged. For suppose we grant (*pace*, e.g., Churchland; Stich) that '[g]etting things right is not just a useful skill. It is a biological imperative' (Dretske, 'Need to Know' 89) – hence, that natural selection would favor *reliable* belief-forming mechanisms; i.e., ones which tend to produce true beliefs. It is still an open question whether SP-type rationality has a monopoly on reliable cognition. For it is not as though the only alternative to SP-rational belief – i.e., belief based on the sorts of 'good rules of inference' singled out by the SP – is SP-*ir*rational belief – that is, belief based on *lousy* versions of those same sorts of rules. Just as genuine an option is SP-*a*rational belief – that is, belief not based on perfectly general, content-neutral rules at all. And, in fact, one main theme in recent work on human reasoning is that humans in fact employ belief-forming processes which, while not resembling anything like SP-rational processes, *are* rather impressively reliable.[9]

5. 'The Rationality Wars' – Psychological and Normative Issues

As an illustration of the latter claim, take those 'biases' the use of which, as we saw previously, some take to be evidence of widespread human irrationality. Tversky and Kahneman are probably the best-known advocates of the 'heuristics and biases' tradition. Yet, while they do argue that such heuristics 'sometimes . . . lead to severe and systematic errors' in reasoning ('Judgment under Uncertainty' 1124), they also observe, more positively, that such heuristics are '[i]n general, quite useful' (1124). In

fact, as Gigerenzer et al. have argued, applied in the right sort of domain – specifically, a domain resembling that in which the heuristics were (allegedly) 'acquired' – using certain 'fast and frugal' heuristics is remarkably reliable (cf. Feldman, 'Rationality' 223ff).

By way of illustration consider, for example, 'the recognition heuristic', which counsels, roughly: *in choosing between two objects, if one is recognized and the other is not, choose the former*. In the case of choosing something to eat, on the assumption that humans have done a reasonably good job of discovering and incorporating edible things into their diet, things we recognize are much more likely to be edible than those we don't recognize. Here, connections, or presumed connections, between various features of the world (between edibility and familiarity, e.g.) underpin and account for the success of the heuristic. Generalizing, 'fast and frugal' heuristics 'exploit the structure of information in the environment to arrive at more adaptively useful outcomes' (Gigerenzer, Todd, and the ABC Research Group 24). In effect, the idea is that *you don't have to be that smart (in the sense of SP-rational) if your beliefs latch on to and reliably track the (non-logical, a-rational) structure of the world*.[10] (For similar ideas, see Brooks; Simon.) Because such heuristics make use of only a fraction of the information available, and because they don't require much at all in the way of 'calculations' on the part of the subject, they are likely to be much less costly than what's required for SP-type rationality. (More on these points below.) Crucially, however, the heuristics' 'frugality' needn't compromise their effectiveness: when applied in the right kind of domain (one resembling what evolutionary psychologists refer to as 'the environment of evolutionary adaptedness') they are rather surprisingly reliable.[11] And yet, they don't look anything like the type of rules apt to be singled out by the SP.

Similar points apply to specific examples of 'reasoning errors' of the sort we canvassed above. If we are already convinced that human psychology contains some formal inference rules which approximate those championed by SP, and that it is the use of these which is importantly responsible for our practical and intellectual success, then it would only be deviations there from which would need explaining – perhaps in terms of various pragmatic principles,[12] the difficulty of retaining and retrieving information, and so on (Evans, 'Theories of Human Reasoning' 86). (Cf. various versions of the 'mental logic' approach to understanding human reasoning – e.g., Braine; Braine and O'Brien; O'Brien; Rips.) This is why, as Evans ('Theories of Human Reasoning') stresses, it is important to see that, among the data to be explained is not merely the apparently disappointing results – namely, our getting the wrong answers on certain reasoning tests; or even that we do so quite routinely and uniformly. Just as theoretically significant are *content effects*: by altering slightly the details of these same reasoning tasks – details which, from a purely formal point of view, are quite irrelevant – performance improves dramatically. Perhaps the best-known

instance of this concerns so-called 'deontic' versions of the Section Task, where, instead of 'abstract' claim, *If a card has a vowel on one side, then it has an odd number on the other side*, subjects are presented with a rule like, *If a person is drinking beer, they must be over 20 years old* (Griggs and Cox). Here, the majority of subjects get the right answer. Why would that be? Various answers have been proposed.

According to Cheng and Holyoak, is it because in confronting such problems, subjects employ a *permission schema* of the form, '*If the precondition is met, then the action may be taken*'. According to Cummins, it is because, due to selective pressures operating within dominance hierarchies such as our ancestors found themselves in, we possess a domain-specific *deontic reasoning module*, though not a set of 'syntactically-driven rules', for addressing questions about permissions, obligations, and so forth. Perhaps the best-known suggestion, however, is Cosmides and Tooby's idea that – again, for selectionist reasons – we've come to possess, in particular, a *cheater-detection module*.[13] There are important and sometimes subtle differences between these views. But for present purposes we can focus on what they share: the commitment to some non-SP means of addressing certain questions which, while perhaps not so reliable when applied outside their proper domain (i.e., one resembling that in which the mechanism, module, or tendency was acquired), are rather impressively reliable when it comes to solving problems that would have been, and continue to be, especially pressing for social beings like ourselves.

Analogous claims have been made in explaining the patterns of performance found across different versions of probabilistic reasoning tasks. For example, returning the example of base-rate neglect in clinical medicine described in the previous Section, when information about base rates is couched in terms of frequencies rather than in straight percentage form, performance increases, sometimes dramatically so (for discussion, see Chase, Hertwig, and Gigerenzer 211). And so too with frequentist versions of the Linda problem (see, e.g., Gigerenzer; Samuels and Stich §3.1): the way in which the problem is framed can dramatically improve performance. Why? Once again (it is suggested), because frequencies, rather than more abstract percentages, would have been, and continues to be, the form in which the relevant information is actually encountered and stored by humans.[14]

Here too, there is plenty of room for disagreement as to the interpretation and assessment of such claims (for which, see, e.g., Bishop). But the salient point, once again, is that we have here the promise of reliable cognition in the absence of the employment of SP-type rules. How this is accomplished is often described in terms of our natural reasoning strategies having certain built-in 'assumptions', or in terms of those strategies being geared towards particular sorts of situations, problem-forms, or environments. Here, an analogy with vision is apt (Chase, Hertwig, and Gigerenzer; Kornblith, 'Laws of Thought'): the visual system has certain

'assumptions' built into it – that certain visible angles correspond with the closer and further edges of an object; that there is a single light source located above (namely, the sun) such that shading is, indirectly, a cue to an object's shape (Ramachandran and Rogers-Ramachandran). Those assumptions may be exploited to misleading effect (the Müller-Lyer illusion, 'countershading' camouflage in animals). But such illusions move no one to a general scepticism about perceptual knowledge. Rather, they are best understood as shedding light on how these generally reliable systems work as well as they do – namely, by exploiting certain features of the environment in which they are employed, thereby dramatically reducing the amount of work that must be done by system(s) themselves.[15] As Kornblith says, the errors to which our belief-forming systems sometimes lead us 'often hide a deeper wisdom' ('Laws of Thought' 911).

In fact, while they tend to emphasize cases in which there is an ostensive failure to reason correctly, Tversky and Kahneman make the very same comparison with vision; and, with the negative emphasis removed, the message is once again much more positive: In vision, blurriness may typically be a good guide to distance, they say, but it won't be such in lousy weather 'when visibility is poor'. So too, they write: various 'heuristic principles' 'reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors' ('Judgment under Uncertainty' 1124).[16]

That the processes singled out by SP do not, after all, constitute the sole route to reliable belief is an under-appreciated fact. Of course, though, if one is working within the confines of SP, even if humans do employ a number of belief-forming methods which are on the whole *reliable*, that is not enough to secure their rationality. Still, a number of objections have been lodged against SP as a correct normative account.

First, recall that epistemic rationality (epistemic reasons, etc.), as intro-duced above (section 2), has to do with the effective pursuit of the distinctively cognitive-epistemic end of true belief. And here, we might pose the following dilemma for SP. Part of the rationale for adopting something like SP was the very natural thought that we must have a grasp of certain formal rules governing logical and probabilistic inference if we are to effectively pursue the truth goal. But if the underlying rationale for our concern with rational belief formation is such a veritistic[17] one, then we're owed some account of why rational belief-forming processes must be restricted to those singled out by SP. On the other hand, if there is supposed to be *something else* which accounts for the value of beliefs formed in the manner SP prescribes, we need to be told what that is. (See Samuels and Stich §8.1, on the 'consequentialist' challenge to 'deontic' theories of rationality.) However, those who seem to be oper-ating with SP in their discussions of rationality, particularly as it regards the reasoning experiments, are characteristically silent on *why* it provides

or constitutes the correct normative standard – they appear to be simply taking it for granted that it does.

Second, as Gilbert Harman (*Change in View*; 'Rationality'; cf. Goldman, *Epistemology and Cognition* 81–9; Samuels, Stich, and Faucher §7.3) has been particularly concerned to stress, we should not confuse inference and good reasoning with implication and consistency. Formal (e.g., logical) rules inform us only of the latter sorts of things. As a general point, though, from the fact that, say, *a & b* implies *b*, or that *a* and *b* are inconsistent, nothing immediately follows about what one should believe. In short, logical rules are one thing, good reasoning, and rules thereof, are another. So, even setting aside the existence of *multiple* formal systems which might be recommended for deductive and/or probabilistic reasoning (Chase, Hertwig, and Gigerenzer 207), when the proponent of SP says that rationality requires reasoning in accordance with principles of reasoning that are based on rules of logic, probability theory and so forth, we need to be told *which* such principles, and *in what way* they are so based.

Lastly, there is the objection that SP-type views promote standards of rationality which it is simply impossible to live up to. So, e.g., that one maintain consistency among one's beliefs seems like the sort of requirement that any SP-type view would impose. But Cherniak (cf. Chase, Hertwig, and Gigerenzer; Goldman, *Philosophical Applications* 20–2), e.g., notes that to do a truth table for just 138 propositions – hence, to check to see if 138 beliefs were consistent – would require a table about 3.5×10^{41} rows long. Similarly, Harman points out that updating the probabilities of propositions, as we'd need to do to employ conditionalization in the case of statistical inference, would likewise lead to 'combinatorial explosion':

> If one is to be prepared for various possible conditionalizations, then for every proposition *P* one wants to update, one must already have assigned probabilities to various conjunctions of *P* together with one or more of the possible evidence propositions and/or their denials. . . . [T]o be prepared for coming to accept or reject any of ten evidence propositions, one would have to record probabilities of over a thousand such conjunctions for each proposition one is interested in updating. (*Change in View* 25–6)

Notice, moreover, that the problem here isn't merely owing to the finitude of humans' cognitive (etc.) resources, as is sometimes suggested. As Kornblith says, 'The problem of combinatorial explosion is one which faces any computational device whatsoever. It is really a feature of the problem itself, rather than a feature peculiar to our, or anything else's, means of solving the problem' ('Laws of Thought' 910). Insofar, then, as we wish to preserve even the *possibility* that humans are rational, given that 'ought' implies 'can', SP seems like a pretty unsatisfactory account of what rationality requires (cf. Botterill and Carruthers 128ff; Goldman, *Philosophical Applications* 22; *Epistemology and Cognition* 279ff; Stein).

Then again, that SP fails as the correct normative account doesn't tell us what to put in its place. But if we are seeking an alternative normative account, we don't have far to look. For instance, the general thrust of much of the discussion of the present section – that such errors of reasoning as we commit don't impugn, and may actually exploit, the *generally reliable* and *actually feasible* means of problem-solving which we employ – suggests that we assay a broadly *reliabilist* conception of rationality. (Such a normative account is also, of course, naturally suggested by the idea that epistemic rationality has to do with the effective pursuit of the distinctively cognitive-epistemic end of true belief.) And a normative account along these lines is precisely what Gerd Gigerenzer and his various collaborators have explicitly advocated. As such theorists see it, SP represents an unacceptable 'opposition between the rational and the psychological', wherein 'rational inference is commonly reduced to logic and probability theory, and psychological explanations are called on when things go wrong' (Gigerenzer and Goldstein 666). SP represents a championing of *coherence* criteria, as opposed to the *correspondence* criteria they favor (Gigerenzer, Todd, and the ABC Research Group 21–2). With the latter, what is stressed is both the *fit* between the 'fast and frugal heuristics' we employ and particular types of situation – the fact that they are suited to work in some environments, not others – and the degree to which their *accuracy* (it is claimed: see note 11) meets or exceeds that of SP-type rules. Putting these ideas together, rationality emerges as essentially a matter of having heuristics which fit the information structure in the local environment, such that using them tends to give you accurate results.[18]

But if SP-type rationality is a wholly unrealistic ideal, it can look as though the foregoing brand of 'ecological rationality' is too sparse to capture everything we want to say about rational belief. For example, there are plenty of situations – many of them quite recent, evolutionarily speaking – in which otherwise ecologically rational heuristics can't be expected to issue in accurate results. As a result, as Grüne-Yanoff says, 'competences adapted to pre-historic circumstances may be of no help in the modern world' (557).[19] Further, there are almost certainly going to be cases in which employing a given heuristic leads the subject to a correct judgment, even though our intuitions balk at crediting her with rational or justified belief. After all, as Gigerenzer says, 'the function of heuristics is not to be coherent' (Gigerenzer et al. 22). And a 'key process in bounded rationality is limited search' (Gigerenzer and Selten 5): the 'frugality' of fast and frugal heuristics consists in precisely such things as not engaging in a Laplaean consideration of all of the information which one might potentially consider. As Bishop puts it:

> *The entire point of the FFH [fast and frugal heuristic] program is that people can, and often should, use very reliable FFHs that ignore lots of evidence and do not properly integrate the evidence they do consider.* (217; emphasis in the original)

For many, such a result would be highly counterintuitive. In short, then, 'correspondence' of the sort that's the hallmark of ecological rationality is, simply in virtue of the general veritistic orientation of most epistemological discussions, a highly desirable, even essential feature of any acceptable theory of rationality. But we *are* capable of certain other, more 'coherence'-oriented forms of cognizing – checking for consistency; deliberately, even ponderously, weighing evidence; reflecting on our belief-forming processes themselves; not to mention, conducting empirical investigations into our own natural belief-forming tendencies so as, perhaps, to ultimately become better thinkers; and so on. And these sorts of more SP-type activities are the sort of thing that many epistemologists have thought to be central to epistemic rationality, and the kind of thing that's required for justified belief and knowledge.

6. Back from the Wars – The Epistemological Debate, Pluralism, and the Search for a Unified Account

In spite of the fact that epistemology has traditionally been carried out without much insight from psychology and the other cognitive sciences, there are some rather striking parallels between the current state of play in 'the Rationality Wars' and various aspects of more traditional epistemological theorizing. Thus, for instance, if the discussion of the previous paragraph sounds familiar, that's not without good reason: it retraces one of the central issues of recent epistemological debate – namely, the nature of justification and, in particular, the debate between 'internalist' and 'externalist' conceptions thereof.

It's difficult to give a characterization of the latter dispute that's both precise and fully satisfactory. (But see, e.g., Steup.) For present purposes, however, the following broad characterization suffices: as noted at the outset, epistemologists are just about unanimous in thinking that whether it is beliefs, their possessors, or the processes by which beliefs are formed that we are assessing, it is a connection with the distinctively cognitive-epistemic end of true belief which is the mark of epistemic rationality (reasons, justification, etc.). What is contentious is what *sort* of truth connection is required for justified belief – in particular, whether it is something of which the subject must be aware or have some 'cognitive grasp' (Bonjour, *Structure of Empirical Knowledge*); whether, 'internal' to her subjective conception of the situation, she must have good reasons or evidence to suppose that such a truth connection obtains;[20] *or* whether the bare fact that her beliefs *are* appropriately connected to the truth – for the reliabilist (e.g., Goldman, 'What is Justified Belief?'), that they are produced/sustained by reliable processes – suffices. Internalists typically maintain that a bare truth connection leaves out such notions as evidence and *responsible* believing – in effect, that it fails to give any central role to reflective epistemic *agents*, but instead treats subjects as though they

are just 'thermometers' registering various items of information, to put it in terms of Armstrong's notorious model.[21]

Externalists, meanwhile, tend to regard it as a *virtue* of their view that it casts humans as being on a continuum with other animals in respect of their cognitive-epistemic connection to the world and enables us to explain why it makes sense to speak, if only figuratively, of certain inanimate objects 'knowing' things (e.g., Dretske, 'Two Conceptions of Knowledge'; 'Need to Know'; 'Précis of Knowledge' 177; Goldman, 'Discrimination and Perceptual Knowledge'; Kornblith, *Knowledge and its Place*). Further, they argue that if those things prized by the internalist have genuine epistemic significance, that is only because and insofar as *they* contribute to the reliability of our belief-forming processes and increase the likely truth of a given belief; but if that's so, it seems arbitrary to insist that justification-conferring factors or processes must always 'somehow be cognitively available to the believer himself, within his cognitive grasp or ken' (Bonjour, in Bonjour and Sosa 24). Finally, as Owens has observed, '[m]uch contemporary work on epistemic justification appears to proceed on the assumption that our cognitive resources are limitless' (194); and externalists often point out that internalistic requirements on justified belief impose unrealistic requirements upon ordinary human subjects – indeed, upon any subjects at all: those demands are simply infeasible and, as a result, internalism in fact tends towards scepticism (see, e.g., Bach, 'Rationale for Reliabilism'; Goldman, 'Internalism Exposed'; cf. Greco 645).

In the epistemological case, resolving such deep disagreement as to the nature of justified belief has proven difficult. So much so, that a number of authors have found it plausible to suppose that the relevant parties are in fact talking past one another; if so, the way to make theoretical progress is to aptly characterize what are, in fact, the multiple explicanda that come into play when we try to understand various epistemological concepts. Thus, to name just a few examples: Kent Bach ('Rationale for Reliabilism') argues that internalists are really concerned with justified (blameless, responsible, etc.) *believers*, not justified (reliable) belief; Keith Lehrer ('Discursive Knowledge') proposes distinguishing between 'primitive knowledge' (roughly, reliably-produced true belief) and 'discursive knowledge' (which requires, i.e., metamentality in the form of critical reflection); Goldman ('Strong and Weak Justification') distinguishes between 'strong' and 'weak' justification; Foley distinguishes between a thoroughly internalistic 'egocentric epistemology' (aimed at answering, from one's subjective point of view, the sort of searching sceptical questions that have traditionally preoccupied epistemologists) and the theory of knowledge (aimed at explaining, from a more third-personal point of view, the conditions in virtue of which a believer is appropriately hooked up to the world); Sosa distinguishes between 'apt' (roughly, reliably-produced) belief and the more internalistic 'justified'

belief, as well as between 'animal knowledge' and that 'reflective knowledge' which incorporates a perspective on one's cognitive faculties and their reliability; and so on.

Perhaps predictably, some of these moves may simply relocate, rather than resolve, the relevant disputes. For instance, Dretske ('Two Conceptions of Knowledge') balks at the very idea that there are, as Lehrer suggests, 'two kinds of knowledge', and Kornblith ('Sosa on Human') reacts similarly to Sosa's distinction between 'animal' and 'reflective' knowledge; and presumably many internalists will insist, *ad* Bach, that justified belief can't be understood apart from the justifiedness of certain 'actions' on the part of believers. So too, while some of these proposals – plausibly, e.g., Sosa's – are quite deliberately fashioned in the hope of constructing an overall unified epistemological theory, others (e.g., Foley's) come with the rider that the prospects for bringing the two (or more) epistemic goods under a single head are rather dim.[22]

Whatever the proper resolution of such issues, however, for present purposes the important point is that many epistemologists have found it plausible that an acceptable theoretical handling of certain central epistemological notions, including justified/rational belief, requires an adoption of *some* kind of '*pluralistic*' view as we might call it. And when we turn back to the psychological literature, we find much the same proposal being made – both with regard to the psychological question of our belief-forming processes themselves, and with regard to the normative question of the nature of rationality.

As to the former, as Samuels and Stich have observed, a movement that has been gaining much attention and many adherents in recent years is the suggestion that human cognition is in fact carried out by two types of reasoning processes. (For a brief overview, see Evans, 'In Two Minds' as well as Stanovich ch. 5; Stanovich and West §6.) While they differ as to details, 'dual-process' theorists generally agree that 'System-1' processes are fast, cheap, parallel, automatic, associative, and largely unconscious, while 'System-2' processes are slower, more resource-reliant, serial, rule-based, and subject to deliberate conscious control. Additionally, System-1 processes are typically thought of as older, shaped by evolution, and geared more towards the solution of problems most pressing to our ancestors. System-2 processes, by contrast, are typically thought to be more recent, more plastic, and much more the product of culture and education. (In light of the latter conjecture, recent data concerning significant individual and cultural differences in reasoning strategies gains a good deal of significance – see, e.g., Nisbett and Norenzayan; Norenzayan et al.; Stanovich; Yates et al., as well as the earlier work of Luria; Scribner; and others.)

Part of the rationale for such a dual-process view, of course, is that it allows us to capture certain salient psychological facts, noted above: for instance, that we seem to have certain non-deliberative, generally reliable

'rules of thumb' which we use in solving certain problems; that these sometimes lead to 'systematic errors'; and yet, that, deliberative creatures that we are, we are able to spot and correct those errors, 'mend our ways', and even develop various explicitly articulated formal rules for solving certain problems. (On the important inhibitory role played by System-2 processes, see, e.g., Sloman, 'Empirical Case'.)

Naturally, since it is itself just a psychological theory, the bearing of the 'dual process' view upon our understanding of rationality remains an open question. But here too suggestions have been made. Evans and Over, e.g., propose a splitting of the notion of rationality which parallels the two systems themselves: 'rationality-1' (or 'personal rationality') involves thinking, speaking, etc., that's generally reliable and efficient for achieving one's goals; while 'rationality-2' (or 'impersonal rationality') involves thinking, etc., on the basis of reasons sanctioned by some normative theory. Along similar lines,[23] Stanovich and West distinguish between 'evolutionary' and 'normative' rationality (see Stanovich; Stanovich and West).

One important issue which arises here, as it did in the epistemological case, is that of the relation between the two (putative) types of rationality. Another is the issue: *in virtue of what* the relevant forms of rationality *are* such. On these questions, there is no settled view. Evans and Over's way of marking the distinction, for example, suggests that rational-1 and rational-2 processes really are distinct types, since they have, on the face of it, different (at least different *proximate*) goals;[24] further, and relatedly, it may seem that rational-1 processes are rational *because reliable*, while the rationality of rational-2 processes is more 'deontic'. As Stanovich and West see it, however, 'evolutionary' and 'normative' rationality may be viewed, not as designating two *types* of rationality, but instead as 'terms for characterizing optimization procedures operating at the subpersonal and personal levels' (661); and 'optimization' here is very naturally read as suggesting a broadly *reliabilist* conception of rationality.[25]

This last point brings out a final point of convergence between the psychological and epistemological literatures. For Stanovich and West's dual-process psychology, along with their pluralistic but ultimately unified account of rationality, closely resembles certain epistemological accounts. For instance, both Pollock and Bach ('Default Reasoning'; 'Rationale for Reliabilism') have emphasized that conscious deliberative thinking is slow, expensive, and ultimately not feasible as the default mode of thinking. Feasible inferences and thinking, then, requires some 'default inference rules' (Bach) or 'quick and inflexible' (Q&I) modules (Pollock) – these play a role analogous to, and may actually *be* or underlie, 'fast and frugal heuristics' or the System-1 mechanisms of dual process theories. Neither view denies, however, that 'we want more than to be automata responding to stimuli with beliefs' (Fumerton 455); and neither denies that there is some good sense in which 'deliberation is the Supreme Court of rationality' (Owens 194). For Pollock,

> the role of ratiocination should be (1) to deal with cases to which Q&I modules do not apply and (2) to monitor and override the output of Q&I modules as necessary. A rational agent can be viewed as a bundle of Q&I modules with ratiocination sitting on top and tweaking the output as necessary. (11)

Similarly for Bach: while a good deal of proper thinking is reliably jumping to conclusions, cognitive competence also involves (reliably) knowing when to think twice.[26] And on neither Pollock's nor Bach's view is there any barrier to thinking that such System-2-type operations are not interestingly influenced or subject to enhancement by experience, culture, education, and so on.

Arguably, whether or not it is explicitly cast in terms of there being 'two systems of reasoning' – much less, whether or not one thinks that there are 'quite literally . . . two minds in one brain' (Evans, 'In Two minds' 458) – it is something like this picture which holds the most promise of providing a satisfactory conception of epistemic rationality. For that picture

(a) places demands upon rational cognition which appear to be perfectly feasible, computationally speaking; it
(b) retains the essential connection between rationality and the distinctive epistemic end of truth; and
(c) while not elevating such things into a free-floating and psychologically unrealistic or otherwise implausible normative standard, it does not downplay or ignore the fact that we are reflective cognitive agents capable of engaging in certain perhaps distinctively human forms of self-conscious, deliberative, 'coherence'-oriented forms of thinking.

In any case, there are other benefits of taking up something like the position described above. First, from the perspective of such a view – again, whether or not it is explicitly cast in terms of there being 'two systems of reasoning' – 'the rationality paradox', as described above (section 4), is really no paradox at all. For there is no longer any pressure to assume that there is a single, uniform set of inferential tendencies, more or less fixed in place, which lead us to err on various reasoning tasks but which are also themselves responsible, not just for vital though mundane tasks like finding food and avoiding cliffs, but also for astounding technical achievements like putting people on the moon.

Second, from this same perspective, it becomes less clear that the question, 'Are humans rational?' *full stop*, is a good one to ask. In many ways, and to varying degrees, we are; in others, not; but the degree to which we are, is perhaps as much a matter of experience, education, and training, as of anything that happened in 'the environment of evolutionary adaptedness'. In view of the fact that, as we've seen, attempts to answer the question of whether humans are rational in some rather more straightforward way are liable not to do justice to the full range of facts

and considerations available to us, that we don't have to give a pat answer to it is surely a good thing.

Short Biography

Patrick Rysiew's primary research interests are in epistemology, including its points of intersection with certain issues in philosophy of language and psychology. Besides rationality, he is currently working on topics including evidence, the epistemology and psychology of testimony and perception, epistemic value, and the pragmatics of knowledge attributions; he has a strong and abiding interest in the philosophy of Thomas Reid. Patrick joined the University of Victoria as Associate Professor of Philosophy in 2006. Before that, he was at the University of British Columbia. His publications include articles in *Noûs*, *Philosophy and Phenomenological Research*, *Australasian Journal of Philosophy*, *The Philosophical Quarterly*, *The Canadian Journal of Philosophy, Analysis*, *The Journal of Scottish Philosophy*, *Consciousness and Cognition*, and *Episteme*. He holds a Ph.D. in Philosophy from the University of Arizona, an M.A. in Philosophy from Dalhousie University, and a B.A. from Simon Fraser University.

Notes

* Correspondence address: Department of Philosophy, The University of Victoria, P.O. Box 3045, Victoria, British Columbia, V8W 3P4, Canada. Email: rysiew@uvic.ca.

[1] A word about terminology. Within epistemology, 'reasoning' is typically reserved for more or less conscious and deliberate ratiocinative thinking. In the relevant psychological literature, however, it tends to used very broadly, such that it is more or less synonymous with 'thinking'. Unless otherwise specified, the discussion here – much of which concerns the latter body of work – assumes this broader sense of the term.

[2] Briefly, Gettier cases involve subjects who satisfy the standard and central three conditions on knowing – justified true belief – but who intuitively shouldn't be counted as knowing, since the desired connection between the justifiedness of their belief and its truth is somehow wanting – there is an important sense in which their (justified) belief is true only 'by accident'.

[3] That it preserves our intuitive anti-scepticism is standardly taken to be a *desideratum* of a satisfactory theory of knowledge; Feldman (*Epistemology* 3), e.g., describes it as 'the first thesis' within what he calls '*the Standard View*'.

[4] Of course – and this will be a running theme in what follows – much depends upon how one conceives of rationality in the first place. Epistemological externalists, e.g., won't regard an assumption of humans' exhibiting *internalist-type* rationality as requisite for epistemological theorizing. So too, externalists tend to stress reliability as the central epistemological concept; and, as we'll see below, it's an open question whether certain familiar experimental findings suggest that we don't satisfy more a reliability-based conception of rationality.

[5] In addition to the individual works cited, the relevant studies are discussed by, i.e., Botterill and Carruthers; Evans, 'In Two Minds'; Samuels and Stich; Stein; Stich. Other important

sources include Gilovich, Griffin, and Kahneman; Kahneman, Slovic, and Tvesky; Nisbett and Ross.

[6] My thanks to an anonymous referee for pressing this point.

[7] In, e.g., Botterill and Carruthers; Kornblith, 'Laws of Thought'; Stanovich; Stein; Stich; Stich and Nisbett. Here, I simply briefly rehearse some of the criticisms such authors present. Readers interested in pursuing these arguments and criticisms in some detail should consult these sources.

[8] See too, e.g., Quine, 'Natural Kinds' 126; Sober 109.

[9] For further critical discussion of the evolutionary argument and the respects in which we should and shouldn't expect to find 'encouragement in Darwin', in Quine's phrase ('Natural Kinds'), see Rysiew.

[10] Here, as they don't affect the present discussion, I am glossing over some of the differences of detail between the heuristics Tversky, Kahneman, and others initially described, and those which Gigerenzer et al. promote. For discussion of the relation between these two theoretical traditions, see Bishop 2006.

[11] Just *how* reliable they are is a matter of some controversy. Again, see Bishop.

[12] While neither does so in service of a defense of the mental logic approach, both Adler and Sperber, Cara, and Girotto appeal to certain general features of rational communication to explain subjects' performance on certain tasks – Adler, on the Linda problem; Sperber, Cara, and Girotto on the Selection Task.

[13] For a brief summary of the evidence favoring the latter proposal in particular, see Botterill and Carruthers 120–3.

[14] Other examples of similarly reliable 'short cuts' in our reasoning strategies are our apparent reliance upon 'the law of small numbers' (see Kornblith, 'Laws of Thought' 910ff) and 'belief' and 'confirmation' biases (see Botterill and Carruthers 125–6).

[15] It is in just these terms that pragmatic principles (see n. 12 above) ought to be understood: they are not mere conveniences, but means of dramatically reducing the search space of possible interpretations of a given utterance, thereby making the intractable tractable.

[16] This is one of those places where, as Samuels, Stich, and Bishop have documented at length, the divisive rhetoric of 'the Rationality Wars' masks what is in fact a deeper, significant agreement.

[17] There are other, non-veritistic forms of 'epistemic consequentialism' which may be proposed – for instance, Stich (ch. 5–6) tries to raise problems for the idea that truth is our epistemic end, and promotes instead a 'pragmatic' theory of cognitive evaluation.

[18] Gigerenzer, Todd, and the ABC Research Group distinguish between ecological, bounded, and social rationality, each of which they take fast and frugal heuristics to exhibit. While distinct, these forms or faces of rationality are not competing but complementary: they all 'look toward the same central goal: to understand human behavior and cognition as it is adapted to specific environments (ecological and social), and to discover the heuristics that guide adaptive behavior' (25; cf. Chase, Hertwig, and Gigerenzer). Here, since it is the 'correspondence'-oriented character of the heuristics, as against the purely coherence-oriented SP-rules, which is most important, I will speak of this view of rationality as the 'ecological' conception.

[19] Grüne-Yanoff contains some very good discussion of different notions of 'bounded rationality', its history, and some of the issues it faces.

[20] '[W]hat the internalist really wants . . . is a view according to which the subjective character of experience, as internally recognizable, provides a genuinely good reason, internally recognizable as such, for thinking that an appropriate modal correlation between such experiences and external situations of a particular sort genuinely obtains' (Bonjour, in Bonjour and Sosa 188; cf. Lehrer, *Theory of Knowledge*).

[21] Of that model, Bonjour says: '. . . the whole idea, central to the western epistemological tradition, of knowledge as essentially the product of reflective, critical, and rational inquiry has seemingly vanished without a trace' ('Can Empirical Knowledge Have a Foundation?' 267). In a similar spirit, Fumerton writes: 'Knowing, or have a justified belief, in the externalist's sense doesn't satisfy our philosophical curiosity, doesn't answer our philosophical questions, because qua philosophers trying to be rational, we want more than to be automata responding to stimuli with beliefs. I would argue that we want *facts*, including facts about which propositions make probable others, before our consciousness' (455).

[22] Another recent example of such a view is provided by Alston, who suggests that there is no single property picked out by 'justification' and that epistemology would be better served by epistemologists' 'dispensing with "justification" '.

[23] On the differences, as they see it, see Stanovich and West 665 n. 11.

[24] For some criticism of this feature of Evans and Over's account, see Oaksford and Chater 294–7.

[25] Analogously, Carruthers argues that System-2 processes are at least partly realized in cycles of operation of System-1. Clearly, understanding the relation between these two putative Systems will be the subject of lively debate for some time to come. So too, there is room for debate as to whether System-2-type processes are achieved via 'mental rules' or 'mental models' (Evans, 'In Two Minds' 457); indeed, there is dispute as to whether the latter two notions are themselves competing (Evans, 'Theories of Human Reasoning' 88).

[26] Cf. Sosa: 'no human blessed with reason has merely animal knowledge of the sort attainable by animals. . . . For reason is always at least a silent partner on the watch for other relevant data, a silent partner whose very silence is a contributing cause of the belief outcome' (240). Cf. too Nozick 74–5.

Works Cited

Adler, Jonathan. 'Abstraction is Uncooperative'. *Journal for the Theory of Social Behavior* 14 (1984): 165–81.

Alston, William. *Beyond 'Justification': Dimensions of Epistemic Evaluation.* Ithaca, NY and London: Cornell UP, 2005.

Armstrong, D. M. 'The Thermometer Theory of Knowledge'. *Belief, Truth and Knowledge.* Cambridge: Cambridge UP, 1973. Reprinted in *Knowledge: Readings in Contemporary Epistemology.* Eds. S. Bernecker and F. Dretske. Oxford: Oxford UP, 2000. 72–85.

Audi, Robert. 'Theoretical Rationality: Its Sources, Structure and Scope'. *The Oxford Handbook of Rationality.* Eds. Alfred R. Mele and Piers Rawling. Oxford: Oxford UP, 2004. 17–44.

Bach, Kent. 'Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice'. *Pacific Philosophical Quarterly* 65 (1984): 37–58.

——. 'A Rationale for Reliabilism'. *The Monist* 68 (1985): 246–63.

Bishop, Michael A. 'Fast and Frugal Heuristics'. *Philosophy Compass* 1.2 (2006): 201–23.

Bonjour, Lawrence. 'Can Empirical Knowledge Have a Foundation?' *American Philosophical Quarterly* 15.1 (1978): 1–13. Reprinted in *Epistemology: An Anthology.* Eds. E. Sosa and J. Kim. Malden, MA: Blackwell, 2000. 261–73.

——. *The Structure of Empirical Knowledge.* Cambridg, MA: Harvard UP, 1985.

—— and Ernest Sosa. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues.* Malden, MA: Blackwell, 2003.

Botterill, George and Peter Carruthers. *The Philosophy of Psychology.* Cambridge: Cambridge UP, 1999.

Braine, Martin D. S. 'Mental Logic and How to Discover It'. *The Logical Foundations of Cognition (Vancouver Studies in Cognitive Science, Vol. 4).* Eds. John MacNamara and Gonzalo E. Reyes. New York, NY: Oxford UP, 1994. 241–63.

—— and David P. O'Brien, eds. *Mental Logic.* Mahwah, NJ: Lawrence Erlbaum Associates, 1998.

Brooks, Rodney A. 'Intelligence without Reason'. Massachusetts Institute of Technology Artificial Intelligence Laboratory, A.I. Memo No. 1293 (1991).

Carruthers, Peter. 'An Architecture for Dual Reasoning'. *In Two Minds: Dual Processes and Beyond.* Eds. J. Evans and K. Frankish. New York, NY: Oxford UP, forthcoming.

Casscells, W. A. Schoenberger and T. B. Graboys. 'Interpretation by Physicians of Clinical Laboratory Results'. *New England Journal of Medicine* 299.2 (1978): 999–1001.

Chase, Valerie M., Ralph Hertwig, and Gerd Gigerenzer. 'Visions of Rationality'. *Trends in Cognitive Sciences* 2.6 (1998): 206–14.

Cheng, P. W. and K. J. Holyoak. 'Pragmatic Reasoning Schemas'. *Cognitive Psychology* 17 (1985): 391–416.

Cherniak, Christopher. *Minimal Rationality.* Cambridge, MA: MIT Press, 1986.

Churchland, Patricia. 'Epistemology in the Age of Neuroscience'. *Journal of Philosophy* 84 (1987): 544–53.

Cohen, L. Jonathan. 'Can Human Rationality be Experimentally Demonstrated?' (Symposium). *The Behavioral and Brain Sciences* 4 (1981): 317–70.

Cohen, Stewart. 'Justification and Truth'. *Philosophical Studies* 46 (1984): 279–95.

Cosmides, Leda and John Tooby. 'Cognitive Adaptations for Social Exchange'. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Eds. J. Barkow, L. Cosmides, and J. Tooby. New York, NY: Oxford UP, 1992. 163–228.

Cruz, Joe and John Pollock. 'The Chimerical Appeal of Epistemic Externalism'. *The Externalist Challenge: Studies in Cognition and Intentionality*. Ed. Richard Schantz. New York, NY: de Gruyter, 2004. 125–41.

Cummins, Denise Delarosa. 'Evidence for the Innateness of Deontic Reasoning'. *Mind & Language* 11.2 (1996): 160–90.

Davidson, Donald. *Inquiries into Truth and Interpretation*. Oxford: Oxford UP, 1984.

Dennett, Daniel. 'True Believers: The Intentional Strategy and Why it Works'. *Scientific Explanation*. Ed. A. F. Heath. Oxford: Oxford UP, 1981. 53–75.

Dretske, Fred. 'The Need to Know'. *Knowledge and Scepticism*. Eds. Marjorie Clay and Keith Lehrer. Boulder, CO: Westview Press, 1989. 89–100.

——. 'Précis of Knowledge and the Flow of Information'. *Naturalizing Epistemology*. Ed. Hilary Kornblith. Cambridge, MA: MIT Press, 1985. 217–37.

——. 'Two Conceptions of Knowledge: Rational vs. Reliable Belief'. *Grazer Philosophiche Studien* 4 (1991): 15–30. Reprinted in *Perception, Knowledge and Belief: Selected Essays*. Cambridge: Cambridge UP, 2000. 80–93.

Evans, Jonathan St B. T. 'In Two Minds: Dual-Process Accounts of Reasoning'. *Trends in Cognitive Sciences* 7.10 (2003): 454–9.

——. 'Theories of Human Reasoning: The Fragmented State of the Art'. *Theory & Psychology* 1.1 (1991): 83–105.

—— and David E. Over. *Rationality and Reasoning*. East Sussex: Psychology Press, 1996.

Feldman, Richard. *Epistemology*. Upper Saddle River, NJ: Prentice-Hall, 2003.

——. 'Rationality, Reliability, and Natural Selection'. *Philosophy of Science* 55 (1988): 218–27.

Fodor, Jerry. 'Three Cheers for Propositional Attitudes'. *RePresentations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge MA: MIT Press, 1981. 100–23.

Foley, Richard. *Working without a Net: A Study of Egocentric Epistemology*. Oxford: Oxford UP, 1993.

Fumerton, Richard. 'The Internalism/Externalism Controversy'. *Philosophical Perspectives, 2, Epistemology*. Ed. James E. Tomberlin. Atascadero, CA: Ridgeview Publishing Company, 1988. 443–59.

Gigerenzer, Gerd. 'How to Make Cognitive Illusions Disappear: Beyond "Heuristics and Biases"'. *European Review of Social Psychology* 2 (1991): 83–115. Reprinted in *Adaptive Thinking: Rationality in the Real World*. New York, NY: Oxford UP, 2000. 241–66.

—— and Daniel G. Goldstein. 'Reasoning the Fast and Frugal Way: Models of Bounded Rationality'. *Psychological Review* 103.4 (1996): 650–69.

—— and Richard Selten. 'Rethinking Rationality'. *Bounded Rationality: The Adaptive Toolbox*. Eds. G. Gigerenzer and R. Selten. Cambridge, MA: MIT Press, 2001. 1–12.

——, Peter M. Todd, and the ABC Research Group. *Simple Heuristics that Make Us Smart*. Oxford: Oxford UP, 1999.

Gilovich, T., D. Griffin, and D. Kahneman, eds. *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York, NY: Cambridge UP, 2002.

Goldman, Alvin. 'Discrimination and Perceptual Knowledge'. *The Journal of Philosophy* 73 (1976): 771–91. Reprinted in Alvin Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press, 1992. 85–103.

——. *Epistemology and Cognition*. Cambridge, MA: Harvard UP, 1986.

——. 'Internalism Exposed'. *The Journal of Philosophy* 96.6 (1999): 271–93.

——. *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press, 1992.

——. *Philosophical Applications of Cognitive Science*. Boulder, CO: Westview Press, 1993.
——. 'Strong and Weak Justification'. *Philosophical Perspectives, 2, Epistemology*. Ed. James E. Tomberlin. Atascadero, CA: Ridgeview Publishing Company, 1988. 51–69. Reprinted in Alvin Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press, 1992. 127–41.
——. 'The Unity of the Epistemic Virtues'. *Virtue Epistemology: Essays on Virtue and Epistemic Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. New York: Oxford UP, 2001. 30–48. Reprinted in *Pathways to Knowledge*. New York: Oxford UP, 2002. 51–70.
——. 'What is Justified Belief?' *Justification and Knowledge*. Ed. George S. Pappas. Boston, MA: D. Reidel, 1979. Reprinted in Alvin Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press, 1992. 105–26.
Greco, John. 'External World Scepticism'. *Philosophy Compass* 2.4 (2007): 625–49.
Griggs, Richard A. and James R. Cox. 'The Elusive Thematic-Materials Effect in Wason's Selection Task'. *British Journal of Psychology* 73.3 (1982): 407–20.
Grüne-Yanoff, Till. 'Bounded Rationality'. *Philosophy Compass* 2.3 (2007): 534–63.
Harman, Gilbert. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press, 1986.
——. 'Practical Aspects of Theoretical Reasoning'. *The Oxford Handbook of Rationality*. Eds. Alfred R. Mele and Piers Rawling. Oxford: Oxford UP, 2004. 45–56.
——. 'Rationality'. *Reasoning, Meaning, and Mind*. Oxford: Oxford UP, 1999. 9–45.
Kahneman, Daniel, Paul Slovic, and Amos Tversky, eds. *Judgment under Uncertainty: Heuristics and Biases*. New York, NY: Cambridge UP, 1982.
Kim, Jaegwon. 'What is "Naturalized Epistemology"?' Ed. James E. Tomberlin. Atascadero, CA: Ridgeview Publishing Company, 1988. 381–405.
Kornblith, Hilary. *Knowledge and Its Place in Nature*. New York, NY: Oxford UP, 2002.
——. 'The Laws of Thought'. *Philosophy and Phenomenological Research* 52.4 (1992): 895–911.
——. 'Sosa on Human and Animal Knowledge'. *Ernest Sosa and His Critics*. Ed. John Greco. Malden, MA: Blackwell, 2004. 126–34.
Lehrer, Keith. 'Discursive Knowledge'. *Philosophy and Phenomenological Research* 60.3 (2000): 637–53.
——. *Theory of Knowledge*. Boulder, CO: Westview Press, 1990.
Lichtenstein, Sarah, Baruch Fischoff, and Lawrence D. Phillips. 'Calibration of Probabilities: The State of the Art to 1980'. *Judgment under Uncertainty: Heuristics and Biases*. Eds. Daniel Kahneman, Paul Slovic, and Amos Tversky. New York, NY: Cambridge UP, 1982. 306–34.
Luria, A. R. *Cognitive Development: Its Cultural and Social Foundations*. Cambridge, MA: Harvard UP, 1976.
Mele, Alfred R. and Piers Rawling, eds. *The Oxford Handbook of Rationality*. Oxford: Oxford UP, 2004.
Nisbett, Richard E. and Eugene Borgida. 'Attribution and the Psychology of Prediction'. *Journal of Personal and Social Psychology* 32.5 (1975): 932–43.
—— and Ara Norenzayan. 'Culture and Cognition'. *Stevens' Handbook of Experimental Psychology, Volume Two: Memory and Cognitive Processes*. Eds. Hal Pashler and Douglas Medin. 3rd ed. New York, NY: John Wiley & Sons, 2002. 561–97.
—— and Lee Ross. *Human Inference: Strategies and Shortcomings*. Englewood Cliffs, NJ: Prentice-Hall, 1980.
Norenzayan, A., E. E. Smith, B. J. Kim, and R. E. Nisbett. 'Cultural Preferences for Formal versus Intuitive Reasoning'; *Cognitive Science* 26 (2002): 653–84.
Nozick, Robert. *The Nature of Rationality*. Princeton, NJ: Princeton UP, 1993.
Oaksford, Mike, and Nick Chater. *Rationality in an Uncertain World*. East Sussex: Psychology Press, 1998.
O'Brien, D. P. 'Mental Logic and Irrationality: We Can Put a Man on the Moon, So Why Can't We Solve Those Logical Reasoning Problems?'. *Rationality: Psychological and Philosophical Perspectives*. Eds. K. I. Manktelow and D. E. Over. London: Routledge, 1993. 110–35.
Owens, David. 'Review of *Cognitive Carpentry*, by John Pollock'. *Mind* 106.421 (1997): 192–5.
Pollock, John L. *Cognitive Carpentry: A Blueprint for How to Build a Person*. Cambridge, MA: MIT Press, 1995.

Quine, W. V. O. 'Natural Kinds'. *Ontological Relativity and Other Essays*. New York, NY and London: Columbia UP, 1969. 114–38.

——. *Word and Object*. Cambridge, MA: The MIT Press, 1960.

Ramachandran, Vilayanurs S., and Diane Rogers-Ramachandran. 'Seeing is Believing'. *Scientific American Mind* 14.1 (2003): 100–1.

Rips, L. J. 'Cognitive Processes in Propositional Reasoning'. *Psychological Review* 90 (1983): 38–71.

Rysiew, Patrick. ' "Encouragement in Darwin"?' *Facta Philosophica* 4.2 (2002): 271–86.

Samuels, Richard and Stephen Stich. 'Rationality and Psychology'. *The Oxford Handbook of Rationality*. Eds. Alfred R. Mele and Piers Rawling. Oxford: Oxford UP, 2004. 279–300.

——, Stephen Stich, and Michael Bishop. 'Ending the Rationality Wars: How to Make Disputes about Human Rationality Disappear'. *Common Sense, Reasoning and Rationality*. Ed. Renee Elio. New York, NY: Oxford UP, 2002. 236–68.

——, Stephen Stich, and Luc Faucher. 'Reason and Rationality'. *Handbook of Epistemology*. Eds. I. Niiniluoto, M. Sintonen, and J. Wolenski. Dordrecht: Kluwer, 2004. 131–79.

Scribner, Sylvia. 'Modes of Thinking and Ways of Speaking: Culture and Logic Reconsidered'. *Thinking*. Ed. P. N. Johnson-Laird. Cambridge: Cambridge UP, 1997. 483–500.

Simon, Herbert. 'Alternative Visions of Rationality'. *Reason in Human Affairs*. Stanford, CA: Stanford UP, 1983. 3–36. Reprinted in *Rationality in Action: Contemporary Approaches*. Ed. Paul K. Moser. Cambridge: Cambridge UP, 1990. 189–204.

Sloman, Steven A. 'The Empirical Case for Two Systems of Reasoning'. *Psychological Bulletin* 119.1 (1996): 3–22.

——. 'Two Systems of Reasoning'. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Eds. T. Gilovich, D. Griffin, and D. Kahneman. New York, NY: Cambridge UP, 2002. 379–96.

Sober, Eliot. 'The Evolution of Rationality'. *Synthese* 46 (1981): 95–120.

Sosa, Ernest. 'Knowledge and Intellectual Virtue'. *The Monist* 68 (1985): 224–45. Reprinted in *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge UP, 1991. 225–44.

Sperber, Dan, Francesco Cara, and Vittorio Girotto. 'Relevance Theory Explains the Selection Task'. *Cognition* 57 (1995): 31–95.

Stanovich, Keith E. *Who is Rational? Studies of Individual Differences in Reasoning*. Mahwah, NJ: Erlbaum, 1999.

—— and Richard West. 'Individual Differences in Reasoning: Implications for the Rationality Debate?'. *Behavioural and Brain Sciences* 23.5 (2000): 645–65.

Stein, Edward. *Without Good Reason*. Oxford: Clarendon Press, 1996.

Steup, Matthias. 'The Analysis of Knowledge'. *The Stanford Encyclopedia of Philosophy (Spring 2006 Edition)*. Ed. Edward N. Zalta. 2006 <http://plato.stanford.edu/archives/spr2006/entries/knowledge-analysis/>.

Stich, Stephen. *The Fragmentation of Reason*. Cambridge, MA: MIT Press, 1990.

—— and Richard Nisbett. 'Justification and the Psychology of Human Reasoning'. *Philosophy of Science* 47.2 (1980): 188–202.

Tomberlin, James E., ed. *Philosophical Perspectives, 2, Epistemology*. Atascadero, CA: Ridgeview Publishing Company, 1988.

Tversky, Amos, and Daniel Kahneman. 'Judgments of and by Representativeness'. *Judgment under Uncertainty: Heuristics and Biases*. Eds. Daniel Kahneman, Paul Slovic, and Amos Tversky. New York, NY: Cambridge UP, 1982. 84–98.

——. 'Judgment under Uncertainty: Heuristics and Biases'. *Science* 185 (1974): 1124–31. Reprinted in *Judgment under Uncertainty: Heuristics and Biases*. Eds. Daniel Kahneman, Paul Slovic, and Amos Tversky. New York, NY: Cambridge UP, 1982. 3–20.

Wason, P. C. 'Reasoning about a Rule'. *Quarterly Journal of Experimental Psychology* 20 (1968): 273–81.

Yates, J. Frank, et al. 'Probability Judgment across Cultures'. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Eds. T. Gilovich, D. Griffin, and D. Kahneman. New York, NY: Cambridge UP, 2002. 271–91.