

Epistemology and the Psychology of Human Judgement, by Michael A. Bishop and J. D. Trout. Oxford University Press, 2005. Pp. xii + 205. P/b £15.50, H/b £45.00.

Fred Dretske began his review of my book, *The Fragmentation of Reason*, with the warning that it would 'get the adrenalin pumping' if you are a fan of epistemology in the analytic tradition (Dretske 1992). Well, if my book got the adrenalin pumping, this one will make your blood boil. Bishop and Trout (B&T) adopt the label 'Standard Analytic Epistemology (SAE)' for 'a contingently clustered class of methods and theses that have dominated English-speaking epistemology for much of the past century' (p. 8), and they make a spirited case for the view that SAE should be abandoned; it's just not worth doing. According to B&T, 'the main problem with SAE is methodological: its goals and methods are beyond repair' (p. 22). For them, the primary goal of an epistemology worth having is *prescriptive*; it should tell us how to go about the business of reasoning. They are 'driven by a vision of what epistemology could be—normatively reason guiding and genuinely capable of benefiting the world' (p. 7). For the most part, they maintain, SAE does not even *try* to guide our reasoning or to make the world a better place, and if it did, the methods it embraces would all but guarantee failure. Those methods 'are suited to the task of providing an account of the considered epistemic judgements of (mostly) well-off Westerners with PhDs in Philosophy',—'a group of idiosyncratic people who have been trained to use highly specialized epistemic concepts and patterns of thought' some of which ordinary folk would find 'strange, foreign, or unfamiliar' (p. 107). And, they ask with good reason, why should *that* be of any help in figuring out how we should go about the business of reasoning, or in making the world a better place? 'It doesn't matter how deeply philosophers may have considered or refined their epistemic judgements. We still need to know what's so great about philosophers' considered epistemic judgements' (p. 112). Though arguments in this vicinity have been around for some time, B&T make the case with clarity and verve.

While B&T's assault on SAE follows a familiar battle-plan, their proposal for how to go about building an epistemological theory worth having is boldly original and unapologetically ambitious. Their starting point is the observation that, while SAE does not help anyone reason better, there are branches of psychology, dubbed 'Ameliorative Psychology' by B&T, that do. The older and larger component of Ameliorative Psychology is work in the tradition of Paul Meehl, whose classic 1954 book, *Clinical Versus Statistical Prediction* showed that simple, easy to use, statistical prediction rules (SPRs) make more reliable predictions than human experts in a variety of domains. In the intervening fifty years, there have been hundreds of additional studies, and just about all of them have the same result. When a human expert is pitted against a simple SPR, the human expert loses (Grove & Meehl 1996). Another branch of B&T's Ameliorative Psychology is the work of Gerd Gigerenzer and his colleagues, who have shown 'how people charged with making high-stakes diagnoses (e.g. about cancer or HIV) can

improve their reasoning' (p. 15)—indeed improve it dramatically—by reformulating problems about probabilities as problems about frequencies. This work, B&T believe, leads to 'recommendations that are bluntly normative: they tell us how we *ought* to reason about certain sorts of problems'(p. 12).

One obvious question to ask at this point is how, exactly, the empirical findings of Ameliorative Psychologists lead to *normative* recommendations. This question is made all the more urgent by B&T's valuable and informative overview of Ameliorative Psychology, in chapter two, which goes out of its way to emphasize just how dramatically *counter-intuitive* some of these recommendations can be. As best I can tell, B&T propose to get from experimental results to normative recommendations by appealing to what they call 'the Aristotelian Principle' which 'says simply that *in the long run, poor reasoning tends to lead to worse outcomes than good reasoning*' (p. 20, italics in the original). In some cases at least, it is overwhelmingly clear that using the reasoning strategies recommended by Ameliorative Psychology will lead to vastly better outcomes than relying on the judgements or predictions of experts. (For an impassioned chronicle of the harm caused when people rely on experts rather than SPRs, see Dawes 1994.) But what about the Aristotelian Principle itself? What reason do we have for accepting it? Some philosophers might be tempted to suggest that the Principle is a conceptual truth which follows from the meaning of 'good reasoning' or reflects philosophers' considered judgements about good reasoning. But since B&T take a dim view of analysing concepts and systematizing philosophers' considered judgements, this is not a strategy that appeals to them. Rather, they maintain, the Aristotelian Principle 'is an *empirical, probabilistic claim*' (p. 20, emphasis added). And here, I confess, they simply leave me puzzled. For if the Principle is an empirical claim that we have any reason to believe, then presumably there must be at least rough and ready ways of determining both when an outcome is good or bad and when the reasoning that leads to it is good or bad. On the outcome side, B&T rely, for the most part, on our intuitive judgements, and I have no problem with this. A reasoning strategy that leads doctors to misdiagnose treatable diseases does indeed produce worse outcomes than a reasoning strategy that leads them to diagnose those diseases more accurately. But how are we to determine whether that reasoning strategy is a *good* one? Not by appeal to its outcome, of course, since that would be legitimate only if we had some reason to accept the Aristotelian Principle. And not by appealing to our intuitions about what counts as a good reasoning strategy, since many of the strategies that B&T recommend audaciously flout those intuitions. So I'm flummoxed. And perhaps B&T are too. Though they never quite address the question head on, in the passage that comes closest they write: '[A]s a practical matter, we contend that any psychologically healthy, reflective person who has chosen to spend their life doing epistemology must accept the Aristotelian Principle' (p. 20). If I am reading them right, B&T are saying that you better accept the Aristotelian Principle because if you don't you're a sick puppy!

Advocating the normative recommendations of Ameliorative Psychology is just the beginning of B&T's project. One of the main goals of the book is to 'articulate the epistemological framework that guides these recommendations' (p. 54). What are the normative epistemic assumptions of Ameliorative Psychologists? Or, to put the question in a slightly different way, what are the features of reasoning strategies that lead Ameliorative Psychologists to recommend them? 'By looking at some of the successes and failures of Ameliorative Psychology', they argue, 'we can identify three factors that tend to contribute to the quality of a reasoning strategy. The epistemic quality of a reasoning strategy is a function of its reliability on a wide range of problems; the strategy's tractability (that is, how difficult it is to employ); and the significance of the problems it is meant to tackle' (pp. 54–5). So the epistemological theory that B&T advocate, which they call 'Strategic Reliabilism', maintains that 'epistemic excellence involves the efficient allocation of cognitive resources to robustly reliable reasoning strategies applied to significant problems' (p. 71). Of course, all of these factors require a fair amount of unpacking, and that unpacking (in chapters three to six) leads to some of the most interesting and original material in the book. It includes detailed discussions of issues which rarely find their way into books on epistemology, like the costs and benefits of adopting a new reasoning strategy, and the challenge posed by the need to assess the significance of problems that a reasoning strategy might address. There is much more in these provocative chapters than I can comment on in a brief review. Whether or not one agrees with the details—and there are many I am inclined to challenge—B&T deserve great credit for raising a host of intriguing questions that the hermetic SAE tradition has largely ignored. If the book gets the attention it deserves, some of these questions will open new and productive areas of interdisciplinary research in which both philosophers and social scientists can play an important role.

There are also some less momentous delights in this exceptionally well-written (and footnote free!) volume. One of my favourites is the insightful critique of Gerd Gigerenzer's attempt to undermine Kahneman and Tversky's conclusions about people's poor performance on probabilistic reasoning tasks by appealing to a frequentist interpretation of probability (pp. 123–7). But best of all is B&T's justifiably outraged assault on Jonathan Cohen's ill-advised attempt to argue that base rate neglect is not an error in medical decision making. Pulling no punches, B&T show that following Cohen's advice would inevitably lead to many unnecessary deaths (pp. 127–33).

Let me end by briefly raising a concern about the overall structure of B&T's project. There is prima-facie tension between B&T's Aristotelian Principle, which links good reasoning with good outcomes, and their Strategic Reliabilism, which links epistemic excellence to 'robustly reliable' reasoning strategies—that is strategies that lead to *true* beliefs. What creates the tension is the familiar observation that, in some very significant situations, having

false beliefs leads to better outcomes than having true beliefs. Though examples are legion, perhaps the best known comes from the work of Shelley Taylor and her colleagues who have shown that ‘positive illusions’ and ‘unrealistic optimism’ in patients with HIV leads to both better psychological coping and slower progression of the infection (Taylor 1989; Armor & Taylor 2002). To put the matter simply, if you have false beliefs you live longer and have a higher quality of life. Other investigators have found similar results in patients with heart disease. This suggests that in trying to extract insights about ‘epistemic excellence’ from Ameliorative Psychology in the Meehl and Gigerenzer traditions, B&T have too narrow a focus. If they take the Aristotelian Principle seriously, then, at least in some domains, good reasoning will be robustly *unreliable*.

Department of Philosophy
Center for Cognitive Science
Rutgers University
doi:10.1093/mind/fzl390

STEPHEN STICH

References

- Armor, D. A. and S. E. Taylor 2002: ‘When predictions fail: The dilemma of unrealistic optimism’, in Gilovich, Griffin and Kahneman, pp. 334–47.
- Dawes, R. 1994: *House of Cards*. New York: Free Press.
- Dretske, F. 1992: ‘The Fragility of Reason’. *Dialogue*, 31, pp. 311–20.
- Gilovich, Thomas, D. Griffin and D. Kahneman (eds) 2002: *Heuristics and Biases: The Psychology of Intuitive Judgement*. New York: Cambridge University Press.
- Grove, W. M. and P. E. Meehl 1996: ‘Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy’. *Psychology, Public Policy and the Law*, 2, pp. 293–323.
- Meehl, P. E. 1954: *Clinical versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*. Minneapolis: University of Minnesota Press.
- Stich, S. 1990: *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: Bradford Books/MIT Press.
- Taylor, S. E. 1989: *Positive Illusions: Creative Self-Deception and the Healthy Mind*. New York: Basic Books.