

Anyone who is suspicious of the favouring relation will be suspicious of Dancy's accounts. The favouring relation is in the general account, so all forms of reasoning involve it. Thus any problem with it infects all the specific accounts.

More seriously, Dancy has not said enough to convincingly argue that his view is neo-Aristotelian. To do this, he must convincingly argue for Part 2 of his account of practical reasoning. It seems to me that Dancy ends up denying Part 2 of this account with his introduction of a practical purpose (pp. 122–4). Further, where Part 2 of the theoretical account is plausible, Part 2 of the practical account is not. Part 2 of the theoretical account states that, in a non-defective case of the reasoning, one comes to believe in a certain way because the states of affairs represented by one's beliefs favour believing in that way. This is plausible because one believes in those states of affairs and solely thereby is also in the business of believing. Therefore, recognition of what belief-type is favoured is enough to lead to the belief-token. Part 2 of the practical account states in parallel that, in a non-defective case of the reasoning, one comes to *act* in a certain way because the states of affairs represented by one's beliefs favour acting in that way. This is implausible because simply believing in those states of affairs does not entail that one is in the business of *acting* at all. Therefore, recognition of what act-type is favoured is not enough to lead to the act-token. In other words, believing in the theoretical case is more direct than acting in the practical case.

Consequently, even if Dancy's accounts are better than the prevailing ones, he has failed to give a positive defence of neo-Aristotelianism. But also, the above problem signals an advantage for prevailing views of practical reasoning: they seem to have a better explanation of the practicality of practical reasoning. On John Broome's account, for instance, practical reasoning starts with intending, so the reasoner is, automatically, in the business of taking action. It is no surprise, then, that practical reasoning characteristically leads to an act-token. And Dancy can't simply modify his account to state that practical reasoning starts with intending because the states of affairs represented by future-directed intentions do not obtain and thus can't favour anything. Therefore, reasoning in general would not simply amount to tracking the favouring relation, and Dancy would lose much of what makes his view powerful.

Ultimately, this book is an excellent attempt to upset our orthodox ways of thinking about reasoning.

doi:10.1017/S0953820819000487

Kristoffer Ahlstrom-Vij and Jeffrey Dunn (eds.), *Epistemic Consequentialism*

(Oxford: Oxford University Press, 2018), pp. 352. \$77.00.

Nathaniel Sharadin 

The College of New Jersey

There's very little contemporary work in epistemology that isn't in some way related to issues in normative or metanormative epistemology. I'm not sure this really is a new

development; only, epistemologists appear, with few exceptions, to have explicitly embraced this turn toward the evaluative. To my mind, it's all to the good. For, plainly, the epistemic realm is an evaluative realm: the character of Giuliani and company's doxastic lives isn't just largely inaccurate, it's poorly grounded, only haltingly articulable, feckless, disingenuous... I could go on. These ways of agents' doxastic lives might go – inaccurately, fecklessly, and so on – are all ways of agents' doxastic lives going less well than they might otherwise go, from an epistemic point of view.

This book is a collection of essays broadly concerned with what the shape of the epistemic point of view is like. In particular, each essay engages with the idea that the epistemic point of view is in some way a *teleological* or *consequentialist* point of view. What does this mean? That's itself a matter of some debate (the chapters by Littlejohn, Wedgwood, and Driver especially address the question). But, very broadly, it means that the application of epistemically evaluative notions (e.g., justification and knowledge) is grounded in the *instantiation* (Littlejohn, I think, and certainly Wedgwood) or *promotion* (everyone else) of certain values. In a familiar, if potentially misleading, slogan, the (epistemic) good is prior to the (epistemic) right.

Which values, and which goods? Again, there's dispute. But, typically, epistemologists being epistemologists, one – perhaps the only – value is supposed to be *truth*, or its degree-theoretic cousin, *accuracy*. Wherever your focus, either on truth and all-out belief, or on accuracy and credences, you'll find yourself with a range of potentially quite worrying puzzles, interesting applications, and hard-to-settle questions. That's, in the main, the recipe for nine of the book's thirteen essays.

First, pick your ingredients. You'll need one or more doxastic state with which to work, either all-out belief (Snow, Horowitz, Askell) or credences (Horowitz, Carballo, Meacham, Caie, Pettigrew, Joyce, Askell, Dunn). You'll also need a generous helping of some epistemic value or other; truth pairs better with all-out belief (hence, Snow, Horowitz, Askell), whereas accuracy best enhances a dish cooked with credences (hence, Carballo, Horowitz, Meacham, Caie, Pettigrew, Joyce, Askell, Dunn). If you're feeling adventurous, you could try sprinkling in a touch of *something else* (Carballo, perhaps Meacham). Whatever the ingredients, cooked correctly either you'll end up with a familiar dish about which you can express dissatisfaction or (perhaps, given certain additional seasoning) content or you'll end up with some new creation about which you can also express dissatisfaction or content.

The other four essays (Littlejohn, Kornblith, Wedgwood, Driver) are mainly programmatic. In our metaphor – I promise, it's dead after this – they purport to source our ingredients (Kornblith) or to offer culinary roadmaps (Littlejohn, Wedgwood, Driver) rather than specific recipes designed to produce a finished dish. Each are worth reading, both for their contribution to the ongoing debate over what, precisely, epistemic consequentialism is supposed to be (Driver, especially), and for their relationship to the authors' other work in the area (in this respect, especially the Littlejohn, Kornblith, and Wedgwood essays).

Let me make three final organizational comments before we start tasting (really: this time is the last time). First, I found the inclusion of Snow and Carballo's essays in the first section of the book (that is, lumped together with the programmatic essays by Littlejohn, Kornblith, Wedgwood, and Driver) a bit confusing; to my mind, they belong in the applied section and the for-and-against accuracy-first section, respectively. This isn't too important, but it might be helpful for anyone looking to structure their reading of the essays (in a graduate seminar, say) in a thematically consistent way.

Second, there is a quite a bit of formal machinery deployed in quite a few of the essays. By my count, seven of the essays deploy formalisms; five of these are probably unmanageable to a reader without at least a passing familiarity with formal epistemology, especially the probability calculus and broadly Bayesian techniques. This is not meant as a criticism; indeed, to my mind the essay with arguably the most formalism – Joyce’s ‘Accuracy, Ratification, and the Scope of Epistemic Consequentialism’ – contains some of the most interesting insights into the project of epistemic consequentialism. And the value of Joyce’s contribution is only enhanced by the three *hors d’œuvres essays* immediately preceding it; those essays both stand on their own as important contributions to the ongoing debate over the workability of credal consequentialism and help set the table for Joyce’s essay. And each of these, as they say, involves the maths.

The final organization is important. The editors of the book, Ahlstrom-Vij and Dunn – the latter also a contributor – have written an excellent introduction that includes not just an overview of the essays in the book (it has that) but also a kind of guide to the culinary landscape (that one was involuntary): it outlines the shape of epistemic consequentialism as a view and the traditional objections facing it. This is good for two reasons. First, it’s good for readers: I’d recommend reading it before diving into the essays. Second, it’s good for me: because their overview of the included essays is so thorough, I’m relieved of having to provide one myself.

In sum, then, this: the book is worth picking up for anyone even vaguely interested in normative and metanormative debates in epistemology (which, as I’ve said, are many if not most of the ongoing ones), and would serve as an excellent collection to be read in a graduate seminar on the topic. Let me close by highlighting a crucial assumption each of the essays in this book makes that, I think, bears reconsidering.

The thing I think bears reconsidering is somewhat unfamiliar; so let’s start somewhere somewhat less so. According to the orthodox picture of practical rationality, the norms of practical rationality, for example, those governing intention, are invariant across changes in actual agents’ psychologies and environments. Such *invariantist* views about practical rationality differ in what norms they think are the norms of practical rationality. Some invariantists are consequentialists, who think that what practical rationality requires is maximizing (or perhaps satisficing) the practical good. Other invariantists are deontologists, who think that what practical rationality requires is (say) doing what one is obliged to do, and forbearing to do what one is forbidden to do. What all these different invariantist views agree on – consequentialist and deontologist alike – is that whatever the norms of practical rationality are, these norms are the same without regard to changes in feature of particular agents’ context, where that context includes facts both about those agents’ psychologies and about their environment. This is not to say that facts about what the norms of practical rationality require do not vary along with changes in agents’ contexts; of course they do. This is to say that the norms themselves do not so vary. Again, this is a familiar picture.

This picture has recently come under forceful attack by philosophers who think we’d profit from adopting an – as they put it – *ecological* conception of practical rationality (especially Jennifer Morton in, for example, her 2011 ‘Toward an Ecological Theory of the Norms of Practical Deliberation’). According to the ecological picture of practical rationality, the norms of practical rationality depend on details about particular actual agents’ psychological capacities and their environments. Drawing on work from psychology, cognitive science, and (behavioural) economics, proponents of an ecological view of practical rationality argue that, among other things, invariantism delivers the wrong results about what it would be rational for actual human agents with actual human

psychologies to do in a range of environments in which they actually (regularly) find themselves. Ecological rationality, these philosophers claim, does better: since it's a view according to which the norms of practical rationality depend on both the actual cognitive capacities of the agent in question and the actual environment in which that agent finds herself, the norms of practical rationality so understood do better, from the point of view of achieving the aims of practical rationality.

Here's the thing about ecological views of practical rationality. If you're an ecologist about practical rationality, you'll think that an entire swath of arguments between consequentialists, deontologists, and, for that matter, 'teleologists' (who perhaps don't want to be lumped with consequentialists) are pretty confused. For the way those arguments go is that one party will wheel out a case where, given what the agent's psychology or environment is like, it'd plausibly be quite a bad idea indeed for her to engage in consequentialist (or deontological, or teleological) reasoning and then go on to conclude from this fact that the relevant view has somehow, in some way, got things wrong about what the normative structure of practical rationality is like. 'See!', the objector will insist, 'that's why consequentialism/deontology/teleology must be wrong/right/foolish/obviously correct', etc. But from the ecologists' point of view, what these arguments actually do is give us good reason to reconsider our assumption that the norms of practical rationality are *invariant*: what we should instead think is that what norms govern agents' practical reasoning depends on facts about actual agents' psychologies and their environments and (crucially) that these facts are not discoverable from the armchair.

You can see where this is going. The orthodox picture of epistemic rationality, and the one that's as far as I can tell simply assumed by each and every essay in this book, is an *invariantist* picture: it's one according to which, whatever the norms of epistemic rationality are like, those norms do not vary along with changes to agents' psychologies and their environments. So, whether the norms are teleological (Littlejohn, Wedgwood), rule-based (Driver), whether they make room for values such as explanation (Carballo) or focus just on accuracy (Pettigrew and Joyce *for*, Meacham and Caie *against*), whether they entail permissivism about rational requirement (Horowitz), or lead to objectionable epistemic enkrasia (Askell) or epistemic free-riding (Dunn), the norms, as it were, *are what they are*. Thus none of the authors who finds fault with some particular epistemic norm (e.g., maximize expected credal accuracy) in some particular circumstances (e.g., Greaves' examples involving violations of the Principal Principle) draws the moral that in environments like *that*, or for agents with the limitations on their cognitive capacities like *that*, the relevant norm is just different. Instead, what epistemologists are inclined to do is *fix up the norms* in some way; for instance, by arguing that there must be norms other than accuracy-only norms (Meacham, maybe Caie) or that accuracy-only norms are misunderstood (Joyce, Pettigrew). Similarly, on discovering that consequentialist norms (maximize true belief) allow for certain purportedly objectionable trade-offs in certain circumstances (Driver, Littlejohn), or that certain false beliefs are 'adaptive' (Snow), or that consequentialist norms recommend free-riding (Dunn), no one appears to be inclined to say that what they've discovered is a new ecology of epistemic requirement: here, agents aren't required to (say) believe in accord with their evidence; there, agents aren't required to (say) avoid false beliefs. Instead, what epistemologists are inclined to do is *fix up the invariant norms* in some way; for instance, by arguing that they avoid trade-offs because they are not *consequentialist* but *teleological* (Littlejohn, Wedgwood), or that adaptive false beliefs can be responsibly held (Snow) or that free-riding might not be all that bad (Dunn). And so on.

Here's the thing about ecological views of epistemic rationality. If you're an ecologist about epistemic rationality, you'll think that this entire swath of arguments between consequentialists of various stripes, deontologists, and teleologists who don't want to be lumped with consequentialists, is pretty confused. At least, it's confused if the arguments are understood as attempts to limn the nature of a world of invariant normative epistemic requirement. What all these cases give us is good reason to reconsider our claim that the norms of epistemic rationality are in fact invariant: what we should instead think is that what norms govern agents' epistemic lives depends on facts about actual agents' psychologies and their environments and (crucially) that these facts are not discoverable from the armchair. Or anyway, so the story goes. Obviously, there's a lot more to say. But since the invariantist assumption appears to be universally accepted by the essays in this book, it's worth having this alternative, ecological, approach on one's radar while working through them.

doi:10.1017/S0953820819000505