Critical Notice of *Epistemic Consequentialism*

KURT L. SYLVAN

University of Southampton

Introduction

Epistemic Consequentialism represents a shooting-star movement nearing its zenith but already passing its peak of apparent solidarity, with clear fault-lines now emerging in logical space. While the editors push for unity and portray epistemic consequentialism (EC) as a single heading under which various specific views in the book fall, several authors reject the push and defect for categorically distinct enterprises. One aim of this critical notice is to consider the implications of these divisions from an alienated non-consequentialist's perspective, one which stands above the fray of number crunching and tries to get back to normative fundamentals. For what emerges most starkly from reading the volume is that for further progress to be made here, more help might be needed from people who aren't primarily formal epistemologists, and perhaps even from symbol-blind old-schoolers like me.[1] However impressive the technical achievements of epistemic utility theory have been, this volume demonstrates that they have sometimes obscured the philosophical issues pivotal for a debate between something worth calling 'consequentialism' and the non-consequentialist alternatives. While some contributors (Driver, Littlejohn, Wedgwood) explore these issues, and the editors and other contributors (Horowitz and Joyce) see some need for categorical distinctions between the projects executed in their circle, insufficient attention is paid to familiar ideas about the structure of normative theories and the distinction between consequentialist and non-consequentialist views.

The review divides into three main sections. In the first and longest section, I consider what the book contributes to our understanding of the nature of epistemic consequentialism, focusing on the editors' introduction and aspects of pieces by Driver, Joyce, Kornblith, Littlejohn, and Pettigrew. Here I applaud distinctions and efforts in some pieces, push for further distinctions made by non-consequentialists ignored in the volume, and sift out some less helpful suggestions. In §2, I determine what light the EC coalition might shed on particular matters of normative importance; here I look at the pieces by Carballo and Dunn together with some miscellaneous implications noted throughout the volume. In §3, I wrap up by examining new challenges to EC in the pieces from Askell, Horowitz, Meacham, Snow, and Caie. The Conclusion summarizes the lessons I recommend drawing from the book, which I found extremely rewarding despite believing EC even less after reading it.

---

[1]While merit in this corner of philosophy often seems assessed on the basis of technical wizardry, it is worth remembering that some of the most celebrated contributions to the area (normative ethics) that it seeks to imitate came from the self-identified symbol-blind, such as Derek Parfit (see MacFarquhar (2011)).

1 What Is EC?

What is this volume supposed to be about? The editors have one answer, while some authors have importantly different answers. According to the characterization with which Ahlstrom-Vij and Dunn begin,[2] a view counts as EC if it proposes that 'epistemic rightness' should be 'understood in terms of conduciveness to epistemic goodness', where epistemic rightness is 'denoted by terms such as 'justification' or 'rationality'' (2). This definition is good in several ways. Any consequentialist view must 'understand' (or better: *explain*) some target normative property in terms of conduciveness to goodness. Hence, it is not enough for a view to qualify as EC-ist on their definition for it just to give a *value-based* account of a target normative property, or even to understand the target property in terms of its being a *means* to goodness in a broad sense that includes constitutive means (an option the editors instead call 'teleological' (3)). I like these implications, since it is important to remember that there are value-based non-consequentialist views, including both some kinds of virtue ethics and Wood (1999)'s and Herman (1993)'s interpretations of Kant.

Unfortunately, though, there are still views which end up counting as consequentialist given their definition that don't seem consequentialist to me. Here Ahlstrom-Vij and Dunn are willing to say that any view which seeks to explain rightness in terms of the goodness of *only some restricted consequences* of some target evaluandum (e.g., an act, rule, method...) is consequentialist. But as they themselves note elsewhere (see Ahlstrom-Vij and Dunn (2017)), this definition classifies some views which countenance *agent-centered restrictions* as consequentialist. I find this an unacceptable stretching of terminology. If the subset of consequences that are selected to explain the normative property are selected by a deontological restriction, the resulting view is simply not consequentialist. Thankfully, other characterizations of EC in the volume improve on the editors'. Littlejohn and Wedgwood both emphasize that *total* consequences must be used to explain a normative property in a genuinely consequentialist theory. Horowitz also rightly notes in her conclusion that a view which identifies the good consequences only by appealing to a further norm is not really consequentialist.[3]

A broader problem with both the editors' definition of EC and some others in the volume (Kornblith's, for example) centers on the fact that there are different *levels of explanation* at which a normative property can be 'understood' in terms of consequences. In ethics, Kagan (1992, 1997) distinguished two kinds of normative theory—factoral and foundational—and proposed that consequentialism is a normative theory of the latter kind. As he rightly insisted, it is not enough that a theory suggests that the only factors that are *relevant* to the attainment of a normative status are consequences. The theory must also say that the *most fundamental explanation* of the status appeals just to consequences. For as I've noted elsewhere (Sylvan (2012, 2014, FC-1, FC-2)) and Horowitz also notes in her contribution, there might be a deeper explanation of why these are the solely relevant factors which is old-school deontological, Kantian, or virtue-ethical. More generally, it is possible to explain one normative status in terms of good consequences while *then* understanding goodness in deontological terms, as Ewing (1947) noted and thought mattered for debunking the appeal of consequentialism. Once these points are appreciated, it is harder to believe the editors' claim that 'much of contemporary epistemology' has been marked by 'a tacit endorsement of epis-

---

[2]The editors also offer a 'big tent' definition which doesn't focus just on the rightness-evaluation of belief, but on the evaluation of other things (e.g., institutions, communities) in terms of good consequences. My comments below apply equally well to this definition.

[3]For related observations, see Sylvan (2012, FC-1, FC-2).

temic consequentialism' (2), since reliabilist views then aren't essentially consequentialist. Without this claim, the class of EC-ists should seem no larger than the class of high-octane formal epistemologists.

While Littlejohn and Wedgwood offer better definitions, they also encourage breaking ranks with EC. Littlejohn draws a useful distinction between EC and 'teleological non-consequentialism', and recommends abandoning the former for the latter partly because the former requires unacceptable tradeoffs. Teleological non-consequentialism agrees that the set of norms governing belief 'tells us that the beliefs we ought to have are the beliefs we ought to have because they *perform a function* or *serve an aim*', but is allegedly also non-consequentialist because 'the good that such beliefs realize is not one that calls for promotion and doesn't determine what's right or rational by ranking options in terms of total value' (44). While I accept the need for some such distinction, I am not sure this is quite the right way to draw it. In particular, I wonder whether it is coherent both to claim that epistemic normativity is fundamentally about 'serving an aim' or 'performing a function' and to deny that epistemic value fundamentally calls for promotion.

A goal, aim or objective is precisely something to-be-brought-about. Anything whose value is that of an aim would call fundamentally for promotion, in a suitably general sense. Similarly, a function is something to-be-fulfilled, and anything whose value is merely that of a functional aim would also call fundamentally for promotion. So for the same reason why I argued in Sylvan (2018: §5) that reliabilist virtue epistemology doesn't avoid the most general version of the swamping problem, I think this view will fail to do so. Functional value and value-as-aim are what I there called *production values*. Something cannot have derivative value relative to these values if it doesn't reliably promote them *in situ*. Yet as Littlejohn elsewhere notes (see 2016), there are important derivative epistemic values that don't reliably promote his fundamental value of knowledge in some worlds precisely because they don't reliably promote truth in those worlds.[4] Compare toasters and demon-worlders. A toaster which is good in this world may not be good in a world with radically different laws of nature, precisely because it fails to toast stuff over there. By contrast, the demon-world intuition suggests that a belief which is good in virtue of fitting one's experiences here is still good in virtue of fitting phenomenally indistinguishable experiences elsewhere.

Wedgwood offers a similar distinction that may improve upon Littlejohn's here in some respects, but probably not all. He also recommends distinguishing between EC and something he agrees (with some reservations) can be called 'teleology' (91), but characterizes the EC/teleology distinction differently. While his teleological view explains norms by values, he is keen to emphasize that the explanation may proceed via the fact that the derivative evaluanda *instantiate* the fundamental value rather than *promote* it. This helps, since the relationship between instantiator and instantiated seems different in some possible cases both from the relationship between a function and that which serves it, and from the relationship between an aim and that which, well, promotes it. Unfortunately, however, it is unclear that this difference helps with the foregoing worry. A value which is 'to be instantiated' will not obviously confer any real derivative value on an evaluandum in a world if it has no reliable tendency to instantiate the value in that world.

Another problem is that the full 'value-based' view that Wedgwood recommends includes a component that Scanlon (1998: Ch.2) regarded as a fundamentally mistaken plank of consequentialism. Note that Wedgwood begins by defining the 'value-based' approach as one on which '*rankings* of *states of affairs*...lie at the heart of normativity' (85). The idea that

---

[4]The argument against a knowledge-first view which explains norms by production values is slightly more complicated than I suggest here; see Sylvan (2014: Ch. 5) for the full argument.

value is primarily embodied in states of affairs is one which is inconsistent with the idea that there are values which are fundamentally to be respected as given things, as Anderson (1993), Scanlon and Parfit (2011: Ch. 10) argued. Personhood, for example, is arguably the fundamental moral value, but it hardly follows that states of affairs in which persons exist are the things which fundamentally embody this value. Respecting personhood has nothing to do with producing instantiations of personhood, and everything to do with responding in a way which is fitting to this value, a value which (unlike a state of affairs) lacks duration and cannot be 'brought about'. Personhood has the 'dignity' sort of value, not the 'good stuff' sort of value.

To avoid these problems, I recommend defecting for a wider value-based normative theory which would include values-to-be-respected as potential explainers of norms, like in Wood (1999)'s interpretation of Kant. For while Wedgwood may avoid tradeoff objections, I don't think he avoids deeper worries I've raised in (2018) and (FC-1) for EC. To explain why rationality has not merely subjective significance of the kind Wedgwood (2017) suggests but objective significance, I think we must view it as constituting respect for a more fundamental value which is non-teleological. Hence, while I think Wedgwood is right to move from EC to a more general value-based approach, his own approach (like Littlejohn's) remains fundamentally too close to EC.

A different departure from the editors' definition is offered by Joyce. This departure may represent the most promising way to preserve a normative epistemology close enough to consequentialism to deserve the name, but restricted enough in scope to avoid tradeoff objections. Joyce suggests that a consequentialist framework is applicable 'only when agents are committed to using the credences they choose to do the sorts of things that credences are characteristically supposed to do, e.g., making estimates of truth values or setting prices for bets' (240). Although I suspect that anything commitment-like in the manner of belief will not admit of consequentialist treatment for fundamental reasons, I agree that some mental actions, such as betting that p, may admit of consequentialist treatment. Indeed, as Owens (2003) noted, guesses aim at accuracy in just the sort of way that some actions aim at the goals that they are meant to achieve, and are better and worse just in virtue of their likelihood of hitting the mark of accuracy. Even if Owens is right that beliefs do not aim at accuracy in this way, Joyce seems right that estimates and bets do. He also seems right to include the evaluation of estimation and some betting as part of epistemology. The suggestion here is hence probably the best way to make it plausible that some epistemic norms are consequentialist.

I am doubtful, however, that it follows that all epistemic norms should be viewed as derivative of the sorts of norms that apply to betting that p: a lot of epistemology takes place on Main Street, in the minds of we commoners, and has little to do with the kind of speculation that occurs on Wall Street, or in the minds of hedge fund managers.[5] Whether the norms governing credence are part of hedge-fund epistemology, I am unsure. Although credence resembles belief in being a state, it is not clearly a commitment. For this reason, it may have different norms, perhaps ones closer to the norms governing estimating and betting that p. But I don't know, since I've never known what ordinary mental state the technical notion of credence is meant to capture. I am fine, however, with these epistemologists stipulating into existence a state that is related to betting that p in virtue of being a disposition to do so. And indeed, it seems clear that there is an ordinary mental state which stands to betting that p as dispositional stands to occurrent, whatever one wants to call it. Since, however, it can be wrong to commit to the truth of propositions which it might be

---

[5]I defend a version of this claim which applies to Sosa's new virtue epistemology in Sylvan (2017).

right to treat as true on the betting markets, I think little follows for the belief-formation of the many rather than the few with enough security to make such betting seem like some ordinary thing.

Driver's contribution offers a much more direct analogy between the epistemology of ordinary belief and the consequentialist ethics of action. Her contribution differs from the foregoing theorists' in suggesting that we should abandon some apparent ordinary intuitions about objective epistemic justification and explain them away by appealing to a distinct account of subjective justification, in the manner of Railton (1984)'s 'sophisticated consequentialism'. Intriguingly, she suggests that this is the right response to trade-off objections. According to Driver, fundamental epistemic norms do sanction tradeoffs, but these are not the norms that we should use as decision-procedures. It was refreshing to see this view get airtime. It is what I expect is ultimately needed to support a genuine consequentialist analogy between epistemology and ethics (or prudence, as Driver notes is the better analogue). Driver's suggestion is, however, made in an exploratory spirit, without a full defense. More is needed to defend bullet-biting about tradeoffs as a principled move defensible on the basis of a Railtonian split between objective and subjective oughts; Railton's distinction is intuitively appealing when one thinks about the paradox of hedonism, but is there a true epistemic parallel to the intuitions behind the paradox of hedonism? Luckily, however, Pettigrew's contribution to the volume provides a lengthier defense of a kind of error theory that would help Driver, and also provides empirical evidence for this error theory.

The last piece which contains important suggestions about the nature of epistemic consequentialism—Kornblith's—seemed to present the analogy at the wrong level. Kornblith suggests that epistemic consequentialism is a naturalistic theory, which on this account is preferable to the non-consequentialist alternatives. This is a surprising claim, given the rich history of non-naturalist consequentialism—consider Sidgwick, Moore, and Parfit—and the fact that consequentialism is standardly understood as first-order theory, not metaethics. It is also implausible given that Kantianism is a non-consequentialist view with potential constructivist foundations. The claim is all the more surprising coming from Kornblith, who elsewhere (2008) defends a separation of the theory of knowledge from the theory of justification on the grounds that knowledge is not a normative relation but rather a natural, causal, one. I think we should stick with Kornblith's earlier ideas and separate epistemic consequentialism and naturalized epistemology.

I'll end this section with a call to action. In thinking about what the volume tells us about the future of any epistemology that even half-resembles consequentialism in ethics, I was reminded of a quote from André Gorz's *Critique of Economic Reason*: 'Economic rationality has gone as far as it can.... It is as if the crisis of economic rationality were the vacant site of another form of rationality which will give meaning to the whole of the development that precedes it' (94). The crisis of EC which this volume marks is an indication that the program has gone too far under its present banner. But I would expect two new theories of rationality to emerge in its wake, not one. The first will be the epistemic microeconomics of Joyce, which does seem to capture something important about one kind of mental action (for better or worse). The other will be the sort of epistemic non-consequentialism I've recently defended. Either may give meaning to the development that precedes it. But I encourage all who believe that epistemology goes beyond the betting markets to press for the latter to define the shape of future epistemology, not the former.

2 What Good Can EC and its Relatives Do?

Let's henceforth use 'EC+' to refer to the coalition of views documented in the last section, including epistemic teleology, hedge-fund epistemology, and genuine epistemic consequentialism. What good can EC+ do? Much of the volume shows the answer to be: 'A lot, apparently!' I will now turn to these contributions in a positive mood, having discharged some spleen. For I agree that if EC+ represented the whole truth about epistemology, there would be much work for it to do.

This work is mostly done by Carballo, Dunn, Wedgwood, and Pettigrew. Carballo suggests that a non-Veritist version of EC+ could help formalize and explain the intuition that there is more epistemic value in true beliefs about a true scientific theory than true beliefs about blades of grass. I am partly unsure. As I suggested in a discussion of understanding in Sylvan (2018: §4.6), it is far from clear that these evaluative intuitions should be captured by a theory of properly epistemic goodness rather than a theory of goods *simpliciter* which happen to be cognitive (consider Moore's discussion of the value *simpliciter* of knowledge in *Principia Ethica*). Still, I agree that the formal tools of EC+ will shed light on that latter sort of value, which is surely important. The question is just whether this counts as epistemology rather than axiology *simpliciter*.

My only other reservations about Carballo's excellent paper were threefold. Firstly, I didn't think it engaged with the best existing attempts to explain these intuitions within Veritist constraints. The piece does not cite Treanor (2014)'s work, for example, which was designed to answer this criticism of Veritism. It also didn't cite much of the older work on this topic in traditional epistemology. Secondly, the paper overlooked the fact that Veritists can accept non-fundamental but also non-instrumental forms of value of the sort Hurka (2001) uses to explain virtue from a consequentialist perspective. I suspect this kind of derivative value would partly explain Carballo's leading intuitions. Thirdly, the paper might have considered old attempts to do something similar, especially in Lehrer (1974)'s early use of epistemic utility theory to illuminate explanatory coherence. Still, Carballo did beautifully illustrate that work at the interface of the philosophy of science, semantics of questions, and EC+ will shed new light on this old topic.

Dunn's contribution assumes EC+ for the sake of argument and uses it to draw attention to a neglected problem: epistemic freeriding. After noting that earlier work by List and Pettit failed to identify realistic cases of epistemic freeriding, Dunn connects empirical work on group problem-solving, philosophical work on group agency, and the EC framework to identify better cases. The contribution is a fine illustration of the fruitfulness of creatively applying EC+'s tools to applied problems. These tools will clearly enrich the literature on group agency. But it is not clear to me how much this redounds to the credit of EC+, since it was not clear to me why we need EC+ to believe that there are cases of epistemic freeriding. In particular, I wanted to hear more about whether EC+ is needed to illuminate our normative intuitions besides merely serving as a helpful tool for identifying the cases. Only the former kind of promise would suggest EC+ to be preferable to the non-consequentialist alternatives. Still, this was one of the book's best papers.

While Carballo and Dunn offer insight into more applied work that EC+ promises to do, Wedgwood and Pettigrew draw attention to more fundamental work. Amidst considering challenges to EC+, they give thumbnail illustrations of explanatory projects they've pursued elsewhere (in Pettigrew (2015) and Wedgwood (2017)). Wedgwood's paper contains a nice summary of his broadly teleological account of the relationship between rationality and correctness, which he uses in (2017) to address the question of why we should be rational.

Although Pettigrew's contribution is mainly a defense of EC+'s predictions about trade-offs, it also contains suggestions about the normativity of rationality. Pettigrew intriguingly claims (229-230) that if we want epistemic rationality to have normative significance, we must sanction tradeoffs. While these ideas were not the central themes in these papers, they do highlight the relevance of EC+ for big-picture normative questions.

Unfortunately, however, this relevance cuts both ways, as we will see shortly. For as the final pieces I will discuss illustrate, there are other deep questions on which EC+ has questionable implications.

3   Why Reject EC?

The book contains more arguments against EC than for it. Indeed, more than half of the papers raise new problems for EC. Besides the theoretical pieces by Meacham and Caie which are properly classified in the table of contents, this bad news for EC includes two outstanding theoretical pieces by Askell and Horowitz oddly placed in the applied section of the volume, and an applied work by Snow oddly placed in the section meant to cover foundations. Of course, I am not surprised that the volume turns up yet more problems for EC. EC never had the popularity that the editors ascribe it in the Introduction, and is not a promising theory in full generality except for instrumental reasons (e.g., its formal tools help to clarify epistemic free-riding). Still, it is striking that a volume which represents the high watermark of this movement should (in effect) put the mark below sea level.

Let's take stock of these problems and consider what they say about EC's fate. Askell's paper demonstrates with outstanding rigor that EC is inconsistent with enkratic principles barring (a) what she rightly suggests is an implausibly subjectivist account of epistemic value, or (b) rationally required certainty in the true epistemic axiology. While I was persuaded by the paper, it could have been clearer about ways in which EC-ists might evade the problem. This minor shortcoming is related to the paper's incomplete references to the literature on the normativity of rationality: while it cites Broome, it does not cite any of Kolodny's work, for example. It is a familiar heresy in that literature to doubt that enkratic principles are robustly normative, and especially the wide-scope readings that concern Askell. EC-ists could embrace this heresy and take their view to provide a new argument for it. Admittedly, this would be uncomfortable for EC-ists who think their project is to vindicate principles of rationality. But they could take the line of Lord (2018) and Kiesewetter (2017) and argue that narrow-scope requirements of reasons come first, with all wide-scope principles being at best froth on the waves. They could then construe EC as a view about these substantive requirements (which would be closer to consequentialism in ethics anyway).

Horowitz's paper raises problems for EC that may be harder to escape. On the surface, her paper concerns the incompatibility of EC and an appealing version of permissivism. Here some EC-ists might respond by abandoning that version. But Horowitz's paper also develops a deeper theme, especially toward the end; this part reveals a fundamental problem for EC (which I've also emphasized from a different angle in Sylvan (FC-1) and (FC-2)). In particular, Horowitz suggests that there are fundamental disanalogies between decision theory and epistemology which make any attempt to justify EC by appealing to the merits of decision theory duplicitous. Given the disanalogy between actual decision theory and 'epistemic decision theory', EC is only defensible with them help of ad hoc consequentializing. Hence, the rational constraints that EC-ists seek to vindicate might be better explained by a non-consequentialist view. I agree with Horowitz that EC fails because it lacks the explana-

tory promise that is often ascribed to it. But she offers a novel reason for this conclusion that I hadn't anticipated. And if her argument succeeds, it potentially reveals a profound dis-analogy between epistemic and practical rationality that would also threaten the movement toward theorizing the two together. Partly for these reasons, it was my favorite paper.

Snow's paper deepens the existing case against EC from tradeoff cases. It describes the empirical literature on *adaptive misbeliefs* (e.g, positive overestimations of self-worth), and argues that we can identify new case-based objections to EC by thinking about some of these beliefs. Although her key cases are related to Berker's examples, they are less easily challenged and have a different intended moral. Her cases are pervasive throughout ordinary life, so it will be harder to bite the bullet about them in a way EC-ists might otherwise be inclined to do. But Snow also uses them to raise an objection about intellectual responsibility rather than justification. She concedes that it might be justifiable from an epistemic point of view to believe that one is fit enough to fight off a surprise attacker if doing so is the only way to increase one's probability of success in the moment. But she argues that this belief is nonetheless not intellectually responsible, since it would have been possible to take earlier actions which would have improved one's prospects of survival without requiring this (temporarily) false belief (65). Since the belief is *only* justifiable now given this earlier failure, it is not a responsible belief. This point highlights a neglected form of intellectual vice in tradeoff cases.

The papers by Meacham and Caie also deepen the lessons of existing objections to EC and complement themes from Horowitz and Snow, though they are dramatically more formal and less accessible. The key contribution of Caie's paper does seem to rest essentially on its formal intensity. With outstanding technical precision, he provides a far more general argument than has yet been offered for thinking that EC about credal norms is radically at odds with pre-theoretic intuition. While the tension his argument reveals is not different in kind from the one already revealed by Greaves, Berker and Caie himself in other work, he does show that there is no need to appeal directly to cases to show that EC flouts common intuition. This is an important advance.

Although it is similarly formally exquisite, Meacham's paper features a core lesson that requires no formal training to appreciate; indeed, the lesson is about how excessive formalism can obscure elementary normative truths. Meacham argues that EC-ists like Pettigrew have failed to give purely Veritist justifications because they commit a simple explanatory fallacy. He explains this fallacy nicely in one sentence: 'Deriving a norm using an alethic utility function and a decision rule doesn't show that we have a purely alethic justification for that norm' (166). This inference fails, Meacham notes, because the decision rule might be profoundly non-Veritist in character. Hence, the mere fact that it conjoined with a Veritist axiology justifies certain norms does not show that these norms have a Veritist justification. The point here is a basic one about the relationship between axiology and deontology (with a small 'd'). Since axiology and deontology are distinct, one could in principle easily combine, say, a hedonist account of final value with a pluralist account of fundamental duties and thereby justify norms that don't look hedonistic. But it would be absurd to claim that this justification is hedonistic if it essentially relies on the pluralist account of duties. It is revealing that such a simple point undermines Pettigrew's response to the alleged inconsistency of accuracy-first epistemology and evidentialism. Ignoring normative ethics in the belief that decision theory is the true normative authority is a great method for producing error.

Conclusion

Especially considering the work surveyed in the previous section, the strong impression one gets from the volume is that the arguments against EC are more numerous and more fundamental than the early anti-consequentialist literature suggested. Berker (2013a, b) made it appear that the main objection to EC was case-based. But most of the new objections in this volume are not case-based. New responses are required, it seems. In the absence of such responses, EC should have exactly as much appeal as it did before people started misclassifying theories like reliabilism as essentially consequentialist: namely, not much. For there is little evidence that EC was a dominant view in epistemology before recent theorists started using the term 'epistemic consequentialism' pick out theories which weren't seriously intended to resemble mirror consequentialism.[6]

Even if EC's meteoric rise will be followed by meteorite breaking and scattering in the dialectical landscape, a complementary and more positive lesson to the one just mentioned also emerges from the book. EC has manifestly succeeded in stimulating a lot of creative new epistemology. The papers in the volume are great and contain groundbreaking new contributions. Although the fact that EC has produced these results doesn't justify believing it, it does justify reading this volume, which I expect to define the shape of future work in this boomtown of epistemology.

References

Ahlstrom-Vij, K. and Dunn, J. 2017. 'Is Reliabilism a Form of Epistemic Consequentialism?' *American Philosophical Quarterly* 54: 183-194.

Anderson, E. 1993. *Value in Ethics and Economics*. Harvard University Press.

Ewing, A. C. 1947. *The Definition of Good*. New York: Macmillan.

Goldman, A. I. 1986. *Epistemology and Cognition*. Harvard University Press.

Gorz, A. 1988. *The Critique of Economic Reason*. Verso Books.

Herman, B. 1993. *The Practice of Moral Judgment*. Harvard University Press.

Hurka, T. 2001. *Virtue, Vice and Value*. Oxford University Press.

Kagan, S. 1992. 'The Structure of Normative Ethics.' *Philosophical Perspectives* 6: 233-242.

Kagan, S. 1997. *Normative Ethics*. Westview Press.

Kelly, T. 2014. 'Evidence Can Be Permissive' in Steup, M. and Turri, J. (eds.) *Contemporary Debates in Epistemology*. Blackwell.

Kiesewetter, B. 2018. *The Normativity of Rationality*. Oxford University Press.

Kornblith, H. 2008. 'Knowledge Needs No Justification' in Smith, Q. (ed.) *Epistemology: New Essays*. Oxford University Press.

Lehrer, K. 1974. *Knowledge*. Oxford University Press.

Littlejohn, C. 2011. *Justification and the Truth Connection*. Cambridge University Press.

Lord, E. 2018. *The Importance of Being Rational*. Oxford University Press.

MacFarquhar, L. 2011. 'How to be Good.' *The New Yorker*, September 11, 2011 issue.

Owens, D. 2003. 'Does Belief Have an Aim?' *Philosophical Studies* 115: 283-305.

Pettigrew, R. 2015. *Accuracy and the Laws of Credence*. Oxford University Press.

Parfit, D. 2011. *On What Matters (Volume 1)*. Oxford University Press.

Railton, P. 1984. 'Alienation, Consequentialism, and the Demands of Morality.' *Philosophy and Public Affairs* 13: 134-171.

Scanlon, T. M. 1998. *What We Owe to Each Other*. Harvard University Press.

---

[6]Goldman (1986) may seem a counterexample. But in a reading group on epistemic value theory I organized at Rutgers, he memorably told us: 'Although I said it [i.e., he once endorsed EC], I never said it loudly.'

Sylvan, K. 2012. *Respect for the Truth and the Normativity of Epistemic Rationality*. Proposal draft of PhD thesis, available on request.

Sylvan, K. 2014. *On the Normativity of Epistemic Rationality*. PhD Thesis, Rutgers University.

Sylvan, K. 2017. 'Can Performance Epistemology Explain Higher Epistemic Value?' *Synthese*, Online first.

Sylvan, K. 2018. 'Veritism Unswamped.' *Mind* 381-435.

Sylvan, K. Forthcoming-1. 'An Epistemic Non-Consequentialism'. *Philosophical Review*.

Sylvan, K. Forthcoming-2. 'Reliabilism without Epistemic Consequentialism.' *Philosophy and Phenomenological Research*.

Treanor, N. 2014. 'Trivial Truths and the Aim of Inquiry.' *Philosophy and Phenomenological Research* 89: 552-559.

Wedgwood, R. 2017. *The Value of Rationality*. Oxford University Press.

Wood, A. 1999. *Kant's Ethical Thought*. Cambridge University Press.