

Christian Piller

How to Overstretch the Ethics-Epistemology Analogy: Berker's Critique of Epistemic Consequentialism

1. Introduction

Should we punish the innocent? Some consequentialists answer, it depends. It depends on how much good may come from punishing the innocent and on how much bad it may prevent. Anti-consequentialists disagree. They tell us not to punish the innocent. An innocent person has done nothing wrong. So we cannot punish her. The benefits others might receive from such punishment are simply irrelevant.

In this paper, I won't argue for this view. It is, I think, a most reasonable view. The question I will pursue is whether this view, anti-consequentialism in ethics, does teach us anything about epistemology.

Compare the following example. Start with some unreasonable view, like the belief that illness is punishment for one's sin. According to this view, if you have a headache, you suffer for your impure thoughts, and if you have AIDS you suffer for your unnatural inclinations. Should we believe that illness is punishment for one's sins? An epistemic consequentialist may answer, it depends. It depends on how much epistemic good may come from holding this view. And like above, many will disagree with such a consequentialist. Even if believing that illness is punishment for one's sins happens to be, for some reason, an entrance ticket to the discovery of important truths in some other domain, this epistemic benefit regarding other truths does not make holding the view reasonable. Like above, how the belief in question affects our other epistemic endeavours is simply irrelevant for whether it can be reasonably held.

We all agree, I assume, on the two normative facts in play. It is wrong to punish the innocent. And it is unreasonable to believe that illness is punishment for one's sins. Selim Berker has argued that these normative facts receive analogous explanations. The ethical consequentialist goes wrong when he claims that benefits to others can outweigh the harms inflicted on the innocent. As Rawls (1971, 26f.) has put it, such thinking would violate the separateness of persons. (I will say a bit more about how to understand this point later on.) The epistemic consequentialist goes wrong when he claims that epistemic benefits in one domain

justify epistemic harms in another domain. As Berker has put it, epistemic consequentialism violates the separateness of propositions.

In this paper I will argue that the idea to appeal to ‘the separateness of propositions’ is misguided. Although there is a sense in which we should consider each proposition ‘on its own merits’, this has nothing to do with the reasons we find in the ethical realm to treat people ‘as individuals’.

2. An Implausible Starting Point

If we want to draw a parallel between the separateness of persons and the separateness of propositions, we need to ask what is meant by those who demand that we treat people as individuals. It could mean that we need to be able to justify what we do to each affected person individually, i.e. by considering the effects our actions have on each person.¹ Let us assume that this is what anti-consequentialists have in mind when they urge us to treat people as individuals, i.e. separately and not as collective. Their claim is a claim about how we ought to relate to each other and, though, it may go beyond our species and encompass all agents, human or otherwise, it won’t apply to other kinds of things. For example, we don’t need to treat grains of sand or grains of sugar as individuals. Whatever it is that marks out humans as a source of obligations in whatever way – and different ethical theories will focus on very different aspects – we certainly won’t find it in grains of any kind. Neither, I want to suggest, will we find what makes humans objects of legitimate moral concern in propositions. Propositions are nothing like people. They don’t cry out to me or put me in some other way under the obligation to consider them, to believe or disbelieve them, whether separately or jointly. In contrast to people, who can hold me to account, propositions are not the kind of thing that could have any normative hold on me.²

Are we not, one might ask, sometimes obliged to think of or consider this or that? When, for example, I promise my friend Karl to read his paper, I ought to consider his ideas, i.e. the bunch of propositions that he endorses in his paper. In examples like these I ought to consider ideas with the required care because I owe it to Karl. I owe nothing, it seems to me, to the propositions he has asked me to think about.

¹ For the development of this idea see Scanlon (1998).

² For the purposes of this paper, we don’t need to consider whether some normative propositions, namely true ones, are an exception to this general claim. Berker’s analogy is not restricted to either normative propositions or to true ones.

To be fair, Berker's idea that we ought to respect the separateness of propositions doesn't entail that the source of this obligation are propositions. I have explained why such an idea would be implausible not in order to provide a master objection to his view, rather to urge for care when we try to find epistemological analogues of moral principles. In the ethical case, it is a reasonable view that each of us deserves our respect. Every person could raise legitimate complaints about what we do and these complaints matter. In the epistemic case, it is not a reasonable view that each proposition deserves our respect or would have a complaint about how we relate to it. So, initially, there is not much plausibility in such a parallel. Berker fails to consider the initial implausibility of the parallel he invokes. Let us look more closely at his ideas in order to find out whether he is able to undermine our initial scepticism.

3. Berker's Project

I will start my discussion by presenting Berker's general project. Consequentialism, in ethics, is a family of ethical theories to be distinguished from deontology on the one side and virtue theory on the other. A rough and ready first formulation of what distinguishes consequentialism from its alternatives is the following structural claim: the good, whatever it is, determines in some way or other what we ought to do. In other words, the evaluative determines the deontic.³ Here is a simple example of a consequentialist view. Pain is bad, so we ought not to bring it about (neither in ourselves nor in others.) In general, we ought to do what makes the world as good as possible. Consequentialism, in epistemology, is likewise a family of theories about what we ought (or are allowed) to believe. It is called consequentialism because it incorporates the same structural idea. We start with an evaluative claim like that believing truly is a good thing and believing falsely is a bad thing. On this basis we argue that we ought to conduct our epistemic lives in ways such that we believe lots of (important) truths and try not to go wrong too often. As it stands this might sound innocent enough but, according to Berker, it is a devil in disguise and we know already why Berker thinks so: it violates the separateness of propositions.

Berker starts his paper 'Epistemic Teleology and the Separateness of Propositions' (Berker 2013a) by highlighting the common normative character of both

³ I will stick with this rough and ready formulation because here is not the place to debate the nature of consequentialism. For a discussion of different ways to separate consequentialism from its rivals see, for example, Pettit (1997) or Brown (2011).

ethics and epistemology. ‘What should I do?’ and ‘What should I believe?’ are, he claims, analogous questions. He previews his paper as follows.

I want to use this analogy between normative evaluations of actions and normative evaluations of beliefs to tease out, and then argue against, a certain strain of thought that seems to have become an article of faith in much recent epistemological theorizing. According to this strain of thought, what distinguishes epistemic norms from other sorts of norms [...] is that epistemic norms are guided by a distinctive set of ‘epistemic’ or ‘cognitive’ or ‘intellectual’ goals. [...] regardless of what the list of epistemic goals looks like, the guiding idea behind this strain of thought is that all other normative notions in epistemology are ultimately explicable in terms of how well the objects of assessment conduce toward, promote, or otherwise subserve these epistemic goals. (Berker 2013a, 339)

Here Berker talks about epistemic goals, later on he uses the notion of the epistemic good. The idea is the same: contemporary epistemology sees the aim of reaching truth and avoiding falsehood (or, alternatively, the idea that truly believing is a good thing and falsely believing a bad thing) as the basis on which to understand other epistemological notions. This priority of the evaluative over the deontic is what we identified as the defining feature of consequentialism. When it falls, much of contemporary epistemology, Berker suggests, will fall with it. ‘I want to convince you’, he says, ‘that this consequentialist/teleological approach to normative epistemology is positively misguided’ (Berker 2013a, 339) and later on he says that this ‘teleological approach to normative epistemology is overwhelmingly the dominant view’ (Berker 2013a, 351).

Consequentialism allows for trade-offs and some trade-offs are morally suspect. Think of trade-offs that would protect many people from some small harm by imposing a big burden on one person. By adding up the small harms which threaten the many we might cross the threshold set by the big burden. However, according to common-sense morality, it is usually no justification for imposing serious harm on someone that many people receive small benefits (by being protected from a small harm). If the many people were one person and so we would face a choice of imposing a big burden on one person or an even bigger burden on someone else, the lesser size of burden can be decisive. In the situation we face, however, in which there isn’t one person who has to carry the whole sum of all the small harms, imposing the burden on the one to benefit the many is, according to commonly held views, impermissible. A utilitarian who thinks otherwise has made a mistake. He sees the situation as if the many could be treated as one big individual and thus he fails to respect the normative principle which Rawls has called ‘the separateness of persons’.

At first sight, the example with which I started – namely that believing that illness is a punishment would open up truths in a different domain – seems to fit

this bill. We sacrifice one belief for adopting many new ones. Berker's latest paper on this issue (Berker 2015) has 'Cutting up the One to Save the Five in Epistemology' in its title. He thinks that epistemologists of a consequentialist bend are committed to endorsing such sacrifices.

[...] if the consequentialist/teleological approach to normative epistemology were the correct one, we'd expect that the correct theory of what we should believe would also countenance trade-offs between goods – in this case, trade-offs between epistemic goods. However, no one – not even those epistemologists who most willingly embrace the consequentialist/teleological framework – is willing to countenance all such trade-offs in the epistemic case. (Berker 2013a, 340)

Let us get clearer about the kind of argument Berker aims to put forward.

4. Two Ways of Pursuing Berker's Project

One kind of argument – for future reference I call it strategy 1 – would emphasise an unreasonable commitment on the part of the epistemic consequentialist. Everyone agrees that, given the evidence available, it would be unreasonable to believe that illness is punishment (even if it opens doors to new knowledge). However, on this line of argument, one tries to show that, regardless of their official denial, epistemic consequentialists are nevertheless committed to telling us that we ought to believe that illness is punishment on the basis of its epistemic benefits in other domains. This commitment, according to our strategy 1 argument, arises from the very nature of consequentialist thought. If we think about epistemic normativity in terms of how best to achieve our end of having a true picture of the world (at least in respect to the parts we are interested in), then, on these consequentialist lines, we should believe that illness is punishment because, by assumption, doing so would open our minds to new and important insights. We lose at one end, in our understanding of illness, but, by assumption, we gain much more in some other area. Overall our aim of understanding the world would be better achieved by adopting this unusual belief. Epistemic consequentialists want to be on our side, i.e. they want to say that seeing illness as punishment would be ridiculous, but the aim-directed nature of their approach commits them to saying what they do not want to say. This is the direction of what I call strategy 1. It could take its inspiration from familiar attempts to show that rule-utilitarianism has to collapse back into act-utilitarianism, or

that it is a form of rule-fetishism that, thereby, is inconsistent with its outward commitment that only truth matters in the epistemic domain.⁴

Berker does not take this line. In order to understand the other line of argument, which is the one he actually takes and which I call strategy 2, we have to remember that consequentialism is a family of theories – just think about the differences between act-, rule-, and motive-utilitarianism. The epistemic analogue of act-utilitarianism has the implausible consequence of condoning our unreasonable illness-belief. This, for Berker, is why epistemic consequentialists have moved away from it. Berker says that examples of the kind illustrated here by the illness-belief and its epistemic benefits, is a counterexample only ‘to extremely crude versions of epistemic consequentialism’ (Berker 2015, 145). The epistemic consequentialist becomes more and more sophisticated and, thus, is not committed to thinking that the illness-belief is reasonable. Berker’s project is to follow the development of consequentialism by adapting his counterexample to the ever new versions of consequentialism. In all these developments he has to focus on the structural features of epistemic consequentialism that allow the development and construction of counterexamples. In Berker (2013a) it is, as the title of the paper suggests, the violation of the separateness of propositions. In Berker (2013b, m8), he identifies ‘epistemic-value-conducivism’ as the structural feature that makes consequentialism vulnerable to counterexamples. ‘The problem with truth-conducivism’, he says there, ‘is that it is a form of epistemic-value conducivism, and conducivism of any sort is the wrong way to think about epistemic normativity.’ I will clarify the idea of truth-conducivism in a moment. In Berker (2015), he explains his argumentative strategy as follows.

The basic idea is to argue that all interesting forms of epistemic consequentialism condone (or mandate) the epistemic analogue of cutting up one innocent person in order to use the organs to save the lives of five people. The difficult part is figuring out exactly what the epistemic analogue of cutting up the one to save the five consists in. (Berker 2015, 145)

So we have to ask Berker, what is the objectionable structural feature of epistemic consequentialism? He reminds us that in (2013a) he thought it was the violation of the separateness of propositions and in (2013b) it was something more general, like truth-conducivism. He still has not reached a definitive answer. ‘But’, he says, ‘that won’t matter for the purposes of this note [which is to reject Goldman’s views]: all the theories I will be considering flout the separateness of

⁴ See Hooker (2000) for a discussion of such objections. Hooker’s replies, for example his distinction between complying with a rule and accepting a rule, have not yet received much discussion in the epistemological literature.

propositions, so there will be no need to go beyond my initial diagnosis.’ (Berker 2015, 146)

Berker’s analysis of what goes wrong in epistemic consequentialism does not, at this point, look satisfactory. If only the most primitive forms of epistemic consequentialism are vulnerable to our counterexample and this example shows why it is important to respect the separateness of propositions, then more sophisticated forms of epistemic consequentialism, which won’t be shaken by this example, need to fail for a different reason. There need to be a different structural feature of consequentialism which allows us to criticise its more sophisticated epistemic forms. The only principle Berker offers is what he calls truth-conducivism. I will follow his lead and explore this idea.

5. What Is Wrong With Truth-Conducivism?

Take an example of a self-fulfilling belief, i.e. a belief about something that is more likely to become true when so believed. If you wholeheartedly believe that you have the abilities that will bring you success, this very fact of believing makes it more likely that you will succeed. Success is, on this kind of view, often a matter of self-confidence. Berker (2013b, m18) gives an example in which an agent, he calls her Jane, believes that she will recover from an illness when this very believing makes what is believed more likely to be true. If you want to reach a true belief about some issue that concerns you, like whether you will recover, and one option – to believe that you will recover – has a tendency for self-fulfilment, could you adopt this belief on this basis? According to Berker, the epistemic consequentialist for whom only truth matters fundamentally – Berker calls this person a ‘veritist’ – needs to say, yes. But Berker insists that self-fulfilling beliefs are not justified by their potential for self-fulfilment.

Now we are at the heart of the matter. Jane’s belief makes it the case that she is more likely to have something which veritism deems to be epistemically good as an end in itself. But I take it to be obvious that her belief is not thereby epistemically justified. Thus the central problem with epistemic consequentialism is not that it neglects the separateness of propositions. Rather, the central problem is its focus on the promoting or conducting relation. (Berker 2013b, m18)

Let us investigate Berker’s claim. Suppose you know more than I do and I ask you about some issue that concerns me. You say, ‘Listen carefully. If you believe what I am about to tell you, you will, by thus believing, have brought it about that you believe the truth about what concerns you.’ How will we expect me to react? Do I need to subscribe to some dubious doctrine of epistemic conse-

quentialism in order to react correctly and believe what I am told? Realising that by forming a particular belief one would bring it about that one will have answered one's question correctly, amounts more or less to having formed this belief. As this example illustrates, adopting truth-conducivism, i.e. the maxim to bring it about that, by forming a particular belief, one has a true belief regarding a matter one is interested in, cannot, in general, be an objectionable feature of a belief-forming strategy.

What about the more particular cases of self-fulfilling beliefs? We have, I think, to distinguish between two kinds of cases. Sometimes the fact that having a belief tends to make true what is believed is not a salient consideration for assessing the truth of what the belief is about. Recovery beliefs in cases of serious and well understood illnesses are normally of this sort. When I get the news that the deadly virus has entered my bloodstream, my belief that confidence matters for recovery won't be salient. I don't believe that confidence can kill viruses and so it would indeed be irrational for me to believe in my recovery, even if recovery beliefs have in general some tendency of self-fulfilment. At other times, self-fulfilling beliefs are justified partly on the basis of their capacity for self-fulfilment. Beliefs that express one's future-directed intentions are of this latter sort. Some of my plans, I am confident, will be realised. My plan is to be in London tomorrow and I believe I will succeed, i.e. I believe that I will be in London tomorrow. Would I not believe this, I would, in this scenario, not really plan to go to London. Thus I will be in London partly because I reasonably believe that I will be in London. The belief that I will be in London plays a role in bringing about its own truth, and, it seems to me, can reasonably be held partly on this basis.

Berker said that 'now we are at the heart of the matter': what is really objectionable about epistemic consequentialism is that it focusses on bringing about what is epistemically desirable, namely to believe truly. So far I have argued that, in general, there is nothing objectionable about believing something in order to bring it about that one believes the truth. This view is supported by the close connection that holds between believing something and taking it to be true. Berker fails to engage with this general connection. He focusses instead on the special case of self-fulfilling beliefs. I have argued that self-fulfilling beliefs, like beliefs about what one is going to do, are sometimes justified partly on the basis that they are self-fulfilling. Some self-fulfilling beliefs, however, do not seem to be justified. The optimism expressed in one's belief in one's recovery, if not supported by appropriate evidence, will be unjustified even if, by some strange coincidence, it turns out to be true. Does this fact – some self-fulfilling beliefs are unjustified – refute a doctrine that is prevalent in contemporary epistemology? At the very least, more needs to be said on behalf of such a conjecture to capture the complexity of the situation. Why is conduciveness to truth gener-

ally acceptable when in some cases of self-fulfilling beliefs the fact that they are self-fulfilling does not determine their justificatory status?

There is a natural answer to this question. Some recovery beliefs are not justified, even if they turn out to be true. Given the way I see the world, recovery in a case where I have been infected by a deadly virus is, from my epistemic perspective, not up to me. Thus, any belief in the power of my optimism would thus be in tension with my way of seeing the world. Contrast this case with one's confidence that one will achieve what one intends to do. Given the way I see the world – with trains running on time that bring me to London in less than 2 hours – whether I will be in London is, from my perspective, up to me. If I believe that I will be in London, I will be. The self-fulfilling nature of this belief, within the setting of my plan of going to London, justifies my so believing.

I have drawn a distinction between problematic and unproblematic self-fulfilling beliefs and it remains unclear why an epistemic consequentialist, for whom reaching the truth is epistemically most important, could not subscribe to the explanation that I offered of why one's optimism in the problematic case would not be justified: it would be in tension with other things one believes and so doesn't promise that one reaches one's aim of believing truly. I conclude that self-fulfilling beliefs pose no special problem for epistemic consequentialism. Sometimes they are justified and sometimes not. The epistemic consequentialist is able to participate in the explanation of when they are and when they are not justified.

6. Is there an Epistemological Analogue of the Separateness-of-Persons Principle?

Let me probe again and ask what, in Berker's view, the objectionable feature of epistemic consequentialism is. It is, he has told us in his latest publication on this topic (Berker 2015), the epistemic analogue of cutting up the one to save or benefit the many. If you sacrifice one in order to save many, you treat the benefits to the many as if they were benefits to one individual and, thus, you violate the idea that we ought to treat persons as separate. Berker's thinking is driven by the thought that there must be people in epistemology who make an analogous mistake. As our example of the belief that illness is punishment for one's sins has illustrated, trade-offs between epistemic goods seem highly questionable and so there has to be some principle that such trade-offs would violate. I will

approach the aim of finding such an analogue principle by looking more closely at the separateness-of-persons principle.⁵

Note first that the separateness-of-persons objection to act-utilitarianism does not condemn all trade-offs. We should not sacrifice one person for the good of the many – this is true. But if one person would be able to save another from significant harm by making a small sacrifice, then one ought to make this sacrifice. For example, even if busy, one should stop and see whether the person who fell down close to us is ok or, if not, one should call an ambulance. The minor loss of convenience for one person is justified by the prevention of a potentially big loss to the other person. The epistemic analogue would be that we are asked to make sacrifices in one domain in order to prevent an even bigger harm in some other domain. On this picture it would be acceptable to give up one epistemic good for another epistemic good which is more significant. For example, I give up my belief that there is milk in the fridge in order to gain some important philosophical or scientific insight. Note that such a transfer would not violate the separateness of persons: the separateness-of-persons principle forbids some forms of aggregation but allows interpersonal comparability. When it comes to morality, small sacrifices for big individual benefit are, often, required; when it comes to epistemology such sacrifices look as bad as the ones that are supposed to violate the epistemic analogue of the separateness-of-persons principle: no promise of a philosophical or scientific insight justifies my abandoning any belief regarding the content of my fridge. So an analogue to the separateness-of-persons principle that could capture all forbidden epistemic trade-offs would have to look rather different from its ethical cousin. This is my first observation: if the parallel between ethics and epistemology ran in the direction Berker takes it to run, one would expect an analogue for allowable transfers but there is no such thing in epistemology.⁶

Looking back at ethics, we can make a second observation. Within personal goodness sums of benefits can be weighed against big sacrifices. It is only when it comes to interpersonal transfers that ethics sets what its defenders would call limits of decency. In the intrapersonal case, the sum of the small benefits of, for

5 For reasons I explained earlier, I am not entirely satisfied by his ambition to find ‘the analogue of the separateness-of-persons principle’, when he said earlier that ‘only the most primitive form of epistemic consequentialism’ violates it. One way to make progress here would be to understand what he is doing as an instance of strategy 1. This strategy, however, would require additional arguments about the collapse of reliabilism, which is a form of rule-consequentialism, to its act-consequentialist base form, which is not supported by anyone in the literature.

6 It is open to Berker to say that this is simply a point where the analogy breaks down. There are permissible trade-offs in ethics, whereas there are no such trade-offs in epistemology.

example having a slightly cheaper energy bill each month, might justify the big sacrifice of investing in a new heating system. Berker will have to say that this is another point at which the analogy breaks down. Whereas there is a virtue, prudence, that covers trade-offs under the umbrella of what is good for a particular person, it is hard to imagine the epistemic analogue of such intrapersonal trade-offs. It would mean that benefits and harms can be traded-off in respect to one particular belief that *p*. I am not sure how best to illustrate this point. Consider the following case. Should one simply disregard one piece of evidence in order to gain access to many other pieces of evidence? Or rather, would the epistemic state that constitutes my disregarding one piece of evidence be justified by the epistemic benefits of being in this state? In the ethical realm I should normally be willing to give up one good thing for many others of the same value. This, however, does not transfer to the epistemic case.

Berker's analogy, I have argued, breaks down at two points. Neither normatively required interpersonal transfers nor normatively required intrapersonal transfers find a parallel on the side of epistemology. Whereas in the practical domain weighing, comparing and trading-off goods is an essential part of practical thought, epistemology has no room for such kind of reasoning. Non-consequentialist ethics forbids some trade-offs, epistemology, it seems, forbids all. That we find impermissible trade-offs in both domains is, I agree, some parallel. This parallel, however, cannot run very deep: the different attitudes towards trade-offs signify a substantial divide between ethical and epistemological accounts of reasoning and normativity.

7. Why Beliefs are Not Justified by their Epistemic Benefits

Many illnesses are well understood. The evidence is clear: illnesses are not punishment for one's sins. Anyone who believed otherwise would not be justified in so believing, even if, in some imagined scenario, such believing would bring important epistemic benefits. Such epistemic benefits are in general irrelevant for determining whether a belief is justified. Why is this so?

Berker says we would violate an epistemological principle which is related to the ethical separateness-of-persons principle. We have found that the epistemological principle which we are trying to explain has a normative shape different from that of its alleged ethical analogue: it simply says that epistemic benefits do not render a belief justified.

There is, I have already argued at the beginning, not much promise in the idea to let propositions play a role analogous to that which persons play in ethics. The mistake in ethics is to treat the many as one. If the many were indeed one, interpersonal sacrifices would be justified. In epistemology, by contrast, it does not matter whether there is one benefit to believing something or many. Neither many nor one benefit will justify any believing. Thus we have one principle – epistemic benefits do not justify beliefs – but the invoked parallel to ethics does little to offer an explanation of this principle.

Epistemic trade-offs don't justify beliefs. This fact, I want to argue, does not show us anything about the nature of epistemic values. When we bracket the question whether beliefs are justified and look at the ways in which we conduct our lives regarding epistemic matters we realize that trade-offs between epistemic goods are commonplace. Suppose I need to say a long number out aloud in order not to forget it. This activity and the concentration it requires are, in the imagined situation, necessary to retain one epistemic good – knowing the number. The same activity prevents me from achieving other epistemic goods. For example, you want to tell me something, but I can't concentrate on what you are saying because I am constantly repeating the number I want to remember. I have to make a comparative judgement about what is more important: knowing the number or knowing what you are trying to tell me. Once we recognise the structure of this example we realise that epistemic trade-offs are common place and, furthermore, that our epistemic concerns mesh naturally with non-epistemic concerns. Information gathering is part of our lives and its costs, epistemic or otherwise, will have to be weighed against its expected epistemic benefits.

We want to explain the following fact – epistemic benefits of believing do not settle the epistemic status of beliefs. I have just argued that the nature of epistemic values, which is by no means hostile to trade-offs, will offer no such explanation. I have argued earlier that pointing to anti-consequentialism in ethics has not resulted in any useful parallels that could help us in our explanatory task. Neither the normative landscape in which the ethical separateness-of-persons principle is set, nor the starting point – propositions exert some normative force on us to be treated separately – would support Berker's view of an analogous explanation of the respective epistemological and ethical points. Beliefs are not justified by their benefits – this will come out true on any ordinary conception of what may and what may not justify beliefs. Berker told us that his anal-

ysis will undermine much of contemporary epistemology. But who will deny the claim that beliefs are not justified by their benefits?⁷

Though Berker offers no explanation of the fact that beliefs are not justified by their epistemic benefits, an explanation is not hard to come by. Whereas actions can be justified via the benefits they promise to provide (be they epistemic or otherwise), beliefs cannot. In order to find an explanation, we need to focus on what it is to believe and on the intimate connection between belief and truth. A believing is, after all, a taking to be true. What I don't take to be true, I do not believe. Consequently, truth will always be an unnegotiable aspect in our assessment of beliefs. The fact that holding a belief offers epistemic benefits does not speak to its truth. Either we hold the belief, i.e. we take what we believe to be true, or we do not. Factors unrelated to truth, are unable to change this state, which, as I said, is the state of taking something to be true. If I don't think that holding a particular belief would be holding a true belief, I do not have this belief. No promise of benefits would be a consideration of the right kind to make me hold something true, when I take it that such promise does not affect in any way whether such believing would be true.⁸

Note that endorsing the connection between truth and believing does not carry any evaluative commitment with it. Even if we say that a belief is correct if and only if true, such a notion of correctness need not be normative. It would be, without doubt, a further step to argue that we (always) ought to have correct beliefs. The very nature of what believing is shows us that we cannot avoid assessing beliefs in terms of their truth. Thus, considerations unrelated to truth cannot get any grip when it comes to forming beliefs. It does not matter for this explanation of Berker's basic fact whether we say that believing truly is the aim of believing or whether we claim that believing truly is a good thing (though, in many cases, it will be).

Take the following example as an illustration of my point. Suppose I keep everything I own in a box. Sometimes I take out one item, some coins let's say, and put another item in my box. I trade one thing for another. My belief box is not like this. I cannot simply take out one thing and exchange it for another thing. What is in my belief box is in there in virtue of my believing it. As long as I believe it, it cannot be taken out. What are not in the box are the things I do not believe. I cannot simply put them into my belief box because believing is an entry ticket without which nothing will be admitted. This picture, naïve as it

7 Roderick Firth (1981, 12), whom Berker often invokes as an ally, says that '... epistemological belief-utilitarianism is not a position that anyone is likely to advocate.'

8 For a related argument to a similar conclusion, see Shah (2006).

may be, is sufficient to explain Berker's fact. If a belief's further epistemic benefits have no bearing on its truth, they have no bearing on our taking something to be true, which, we have said, is what believing is.

We can make the same point in another way. There is no notion of something being instrumentally true. A belief that leads to other true beliefs is a means to truth but that does not make it true in any sense: instrumental truth is not a kind of truth. Beliefs are, in virtue of their nature, assessable in terms of truth. Their capacity for bringing about truths is, without a notion of instrumental truth, irrelevant for this assessment. We can assess actions, including attempts to forget or attempts to keep something in mind, in terms of their epistemic and other benefits. The same goes for the employment of methods, the following of rules, and the involvement in or the endorsement of practises – all these things can be assessed in terms of how they further our epistemic interests. Beliefs can be assessed in many ways but their assessment in terms of truth is non-negotiable. That is why other forms of assessment will never be able to fully capture epistemic normativity.

8. Conclusion

I said at the beginning that the epistemic consequentialist might tell us that we ought to conduct our epistemic lives in ways such that we believe lots of (important) truths and try not to go wrong too often. I also said that this sounded innocent enough. We now know that it really is innocent. How to conduct one's life, including its epistemic aspects, is, at least to some extent, a matter of weighing goods.

Reading Berker has reminded me of the promise I mentioned earlier to read Karl's paper. Berker talks about 'the spectre that is haunting contemporary epistemology: the spectre of reliabilism' and, in sharp contrast to Karl, he describes his work as an attempt of exorcising this ghost (Berker, 2013b, 363). One of his commentators (Goldman 2015), the author of what we could call 'The Reliabilist Manifesto', has accused him of misunderstanding his opponents in epistemology. Goldman emphasises that adopting or endorsing a reliable method and, on this basis, its deliverances, is not a way of judging beliefs by their benefits. In this paper I have looked at Berker's attempt to learn from non-consequentialism in ethics. I have argued that there is hardly any parallel between the epistemologist's general adversity to allowing epistemic trade-offs for a certain purpose, namely for the purpose of determining the justificatory status of beliefs, and the non-consequentialist's claim that it would lead ethical thought astray if we simply summed up the good and the bad things and disregarded their distri-

bution over individuals and groups. Berker is taken by the fact that when believing one thing leads to believing other true things this, by itself, does not determine the justificatory status of the belief. Berker, I have suggested, is mistaken in thinking that this thought, which I feel is generally agreed upon, has much leverage in the contemporary debate in epistemology. An uncontentious assumption about the nature of belief – believing is a taking to be true – is able to explain this fact. This assumption is available to the epistemic consequentialist like to anyone else. Goldman (2015) looked at Berker’s ideas from the reliabilist’s perspective. I tried to look at the analogy which Berker appealed to from an ethical perspective. In our verdict we agree.

References

- Berker, S. (2013a). Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122(3), 337–393.
- Berker, S. (2013b). The rejection of epistemic consequentialism. *Philosophical Issues*, 23(1), 363–387.
- Berker, S. (2015). Reply to Goldman: Cutting Up the One to Save the Five in Epistemology. *Episteme*, 12 (Special Issue 02), 145–153.
- Brown, C. (2011). Consequentialize This. *Ethics*, 121(4), 749–771.
- Firth, R. (1981). Epistemic merit, intrinsic and instrumental. *Proceedings and Addresses of the American Philosophical Association*, 55, 5–23.
- Goldman, A. I. (2015). Reliabilism, veritism, and epistemic consequentialism. *Episteme*, 12 (Special Issue 02), 131–143.
- Pettit, P. (1997). The consequentialist perspective. In M. Baron, P. Pettit, and M. Slote (eds.), *Three Methods of Ethics*. Blackwell.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press.
- Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.
- Shah, N. (2006). A new argument for evidentialism. *The Philosophical Quarterly*, 56, 481–498.

