

# Epistemic Consequentialism: Haters Gonna Hate

Nathaniel Sharadin  
natesharadin@gmail.com

## **Abstract**

Epistemic consequentialism has been charged with ignoring the epistemic separateness of propositions and with (thereby) allowing trade-offs between propositions. Here, I do two things. First, I investigate the metaphor of the epistemic separateness of propositions. I argue that either (i) the metaphor is meaningfully unpacked in a way that is modeled on the moral separateness of persons, in which case it doesn't support a ban on trade-offs or (ii) it isn't meaningfully unpacked, in which case it *really* doesn't support a ban on trade-offs. Second, I consider the strategy of arguing against the trade-off permitting conception of epistemic rationality that is central to epistemic consequentialism on the basis of our intuitive verdicts about the permissibility of trade-offs in cases. I argue the execution of this strategy suffers a methodological mistake that, once corrected, vitiates the probative value of those intuitive verdicts. Hence the conclusion: the separateness of propositions provides no support for a ban on trade-offs, and an influential independent argument for such a ban is flawed.

**Keywords:** *veritism, epistemic consequentialism, trade-offs, separateness of propositions*

## **1. Introduction**

Something's *final value* is value it has that doesn't depend on its standing in any relation to anything else.<sup>1</sup> Familiarly, a *hedonistic* theory of *practical* value holds that:

1. Pleasure has positive final practical value.
2. Pain has negative final practical value.
3. Nothing else has (positive or negative) final practical value.<sup>2</sup>

Hedonism is naturally paired with the following practical first-order normative theory:

**Practical Outcome-Consequentialism (POC):** An agent *A*'s action  $\phi$  at time *t* is practically correct iff *A*'s  $\phi$ -ing at *t* is included in the practically best outcome available to *A* at *t*.

---

<sup>1</sup> The issues here -- terminological and otherwise -- are fraught. For more, see Bradley (2006), Zimmerman (2001), and Korsgaard (1983). For present purposes, my use is simply stipulative.

<sup>2</sup> See, for perhaps the most famous instance, Mill (1861). For contemporary versions of roughly this sort of hedonism, see Feldman (2004), Bradley (2009), and Heathwood (2006).

Together, hedonism and POC tell us that the practical correctness of an action is a function of the pleasure and pain present in outcomes available to an agent at a time.<sup>3</sup>

Less familiarly, a *veritistic* theory of *epistemic* value holds that:

1. True belief has positive final epistemic value.
2. False belief has negative final epistemic value.
3. Nothing else has (positive or negative) final epistemic value.<sup>4</sup>

Veritism is naturally paired with the following *epistemic* first-order normative::

**Epistemic Outcome-Consequentialism (EOC):** An agent *A*'s belief that *P* at *t* is epistemically correct iff *A*'s belief that *P* at *t* is included in the epistemically best outcome available to *A* at *t*.

Together, veritism and EOC tell us that the epistemic correctness of believing some proposition is a function of the truth and falsity present in the outcomes available to an agent at a time. From here on I'll talk about beliefs being epistemically *rational*, *reasonable*, or *justified* rather than merely *correct*.<sup>5</sup>

Recently, the package of veritism and EOC has come under attack on a number of fronts.<sup>6</sup> Here, I'm interested in one particular criticism: that this package allows for illicit "trade-offs" between propositions. The criticism is simple to state but, as we'll see, it's difficult to pin down. The simple way to state the criticism is via an example. I'll borrow one from Selim Berker, who has recently pressed the case against the pairing of veritism and EOC:

**Scientist:** Suppose I am a scientist seeking to get a grant from a religious organization. Suppose, also, that I am an atheist: I have thought long and hard about whether God exists and have eventually come to the conclusion that He does not. However, I realize that my only chance of receiving funding from the

---

<sup>3</sup> 'Correctness' is a placeholder for other normative terms, e.g., 'obligatory', 'right', etc. These two ideas *don't* tell us plenty, such as how to "weigh up" pleasures and pains, what it means for an outcome to be available to an agent at a time, whether our obligations are to be understood as *actualist* or *possibilist* (c.f. Jackson & Pargetter (1986)), and so on. We can safely ignore these complications.

<sup>4</sup> See Berker (2013a, 2013b, 2015), Goldman (1999, 2015). This is what Pritchard (2010) calls "T-monism", Zagzebski (2004) calls "epistemic value monism" and Ahlstrom-Vij dubs "veritistic value monism" (2013). The 'monisms' here highlight clause (3) in our definition.

<sup>5</sup> Here, I follow the opponents of EOC, especially Berker (2013a). More on this in a moment. As with the pairing of hedonism and POC, there is plenty we can't read off the combination of veritism and EOC. But, again, we can safely ignore these complications. For doubts about the parallel between the two pairings, see Horowitz (forthcoming).

<sup>6</sup> On theoretical grounds, as in Littlejohn (forthcoming, 2015) and Berker (2013a, 2013b, 2015), and on empirical grounds, as in James (2017).

organization is to believe in the existence of God: they only give grants to believers, and I know I am such a bad liar that I won't be able to convince the organization's review board that I believe God exists unless I genuinely do. Finally, I know that, were I to receive the grant, I would use it to further my research, which would allow me to form a large number of new true beliefs and to revise a large number of previously held false beliefs about a variety of matters of great intellectual significance.<sup>7</sup>

According to Berker, "most of us" would agree that forming a belief in God's existence in these circumstances would not be epistemically reasonable, rational, or justified.<sup>8</sup> The problem, according to Berker, is that if we accept both veritism and EOC then we appear to be barred from saying this agreeable thing. This is because, were I to form a belief in God's existence in these circumstances, I'd be making what according to the pairing of EOC and veritism looks like an excellent *trade-off*:

I would be trading off the pursuit of one thing of final epistemic value (true belief and the avoidance of false belief with regard to the proposition <God exists>) in order to greatly facilitate my pursuit of a variety of other things of final epistemic value (true belief and the avoidance of false belief with regard to the propositions that make up the subject matter of my scientific research).<sup>9</sup>

Such trade-offs between *propositions* licensed by EOC and veritism are reminiscent of trade-offs between *persons* licensed by some versions of POC and hedonism: just as the latter view seems to sometimes require us to sacrifice one person's pleasure for the "greater practical good," the former seems to sometimes require us to sacrifice one proposition's truth for the "greater epistemic good."<sup>10</sup> However, says Berker, in the epistemic case such trade-offs are always illicit, and this is because

whereas in the ethical case there is some room for arguing that it is not a normative mistake to ignore the separateness of persons, [...] in the epistemic case there is no wiggle room: *the epistemic separateness of propositions is nonnegotiable*.<sup>11</sup>

But: Is it? I am not so sure. In what follows I have two principal aims. The first is to figure out what, exactly, the "epistemic separateness of propositions" amounts to and, along the way, to figure out why, exactly, philosophers such as Berker think it's nonnegotiable. Since the *epistemic separateness of propositions* is supposed to be an analog to the *moral*

---

<sup>7</sup> Berker (2013a, pp. 363-4). Berker himself borrows this example from Fumerton (2001) who in turn adapts it from one due to Firth (1978).

<sup>8</sup> Berker (2013a, p. 364).

<sup>9</sup> *ibid.*

<sup>10</sup> Berker (2013a, p. 373, p. 375).

<sup>11</sup> Berker (2013a, p. 365).

separateness of *persons*, I'll explore arguments for the nature and negotiability of the former in parallel with arguments over the nature and negotiability of the latter. Armed with a clearer understanding of the separateness of propositions, I'll suggest that *hardcore anti-trade-off-ers* (henceforth *haters*, for short) of Berker's sort face a dilemma. On the one hand, haters can model their understanding of the epistemic separateness of propositions on the moral separateness of persons. But, I'll argue, we have been given no reason to think the separateness of propositions, so understood, is absolutely nonnegotiable. Indeed, as I'll explain, the moral analogue militates in *favor* of negotiability. On the other hand, haters can decline to use the moral analogy. But then we are left with no meaningful content to the idea that propositions are epistemically separate. In that case, we clearly have no reason to accept that propositions are in any interesting (non-metaphorical, non-question-begging) way epistemically separate and so, *a fortiori*, no reason to think their epistemic separateness is absolutely nonnegotiable. My second principal aim is to diagnose why haters might have been led to think that the epistemic separateness of propositions was absolutely non-negotiable, and so that trade-offs are always banned. I'll argue that they've been led to this view via a methodological mistake. Let me make two brief preliminary remarks before getting on with things.

First: as I noted above *all* haters, i.e., all those who think that epistemic trade-offs are always forbidden, are my target. But Berker in particular will be my foil: I'll take him as representative of the class of views that finds something seriously disturbing about the idea that it can be epistemically permissible to trade off one proposition's truth against another's. I focus on Berker for two reasons. First, Berker is clearest, among haters, about what the structure of the argument against trade-offs is supposed to be (this will be important in pursuing my second principal aim in the latter half of the paper). Second, Berker is, as far as I can tell, alone among haters in explicitly appealing to the analogues between the separateness of propositions, on the one hand, and the separateness of persons, on the other. Since I'm going to spend the first half of the paper arguing that this analogy isn't helpful for haters, it might seem I'm loading the die by focusing on Berker in this way (since only he appeals to it). But as a matter of fact I think it's a *good analogy*, in the sense that it gives voice to the relevant sentiment that is behind haters' thinking. If Berker hadn't offered the analogy on the part of haters, I'd happily do so myself. Unfortunately, as I'll argue in the first half of the paper, the analogue doesn't help -- and might well undermine -- haters' position. But that doesn't mean it's not worth exploring in some detail.

The second preliminary remark is this: if I'm successful in achieving my aims, I'll have defended the pairing of EOC and veritism against the charge that it allows for objectionable trade-offs. I'm not the first to offer such a defense. But previous defenses have all, in various ways, attempted to show that haters' understanding of veritism or of EOC are somehow flawed. The idea behind these previous defenses, then, has been that EOC and veritism do not have the implausible result haters ascribe to it: contrary to what haters say, EOC and veritism do not "ignore the separateness of propositions" or allow for the trade-offs to

which haters object.<sup>12</sup> By contrast, my defense of EOC and veritism begins by accepting haters' charge that EOC and veritism sometimes ignores the separateness of propositions and allows for trade-offs. I'm interested in what, precisely, it means to say this and why, precisely, we should care if it's true. What I'm going to argue is that if we make good sense out of the claim that propositions are epistemically separate, then it turns out we shouldn't think such separateness is nonnegotiable; if instead we make a hash out of the idea of propositions' separateness, then their separateness *really* isn't nonnegotiable. Let's get on with it.

## ***2. The Moral Separateness of Persons & The Epistemic Separateness of Beliefs***

The claim that propositions are "separate" is a metaphor: to be meaningful, it must be unpacked. We are told that propositions are "*epistemically* separate," but this turns out to mean that propositions are the sort of thing such that trading one's truth off against another's is always illicit. That's obviously not much help in the present context. Happily, haters do provide some guidance: they point us to the separateness of different things in a different normative domain they think is analogous to the separateness enjoyed by propositions, viz. the *moral* separateness of *persons*. So, let's take them at their word and explore the analogy in the hopes of refining our understanding of the separateness of propositions.<sup>13</sup> What does it mean to say that persons are morally separate?

The claim that persons are morally separate is no less metaphorical. But proponents of the idea in the moral domain have said something somewhat more concrete about what they mean by this idea. Tellingly, moral philosophers regularly appeal to the separateness of persons in the course of arguing against versions of the pairing noted above between hedonism and POC. Since epistemologists such as Berker appeal to the analogous idea in the course of arguing against the analogous pairing of veritism and EOC, it seems we might be on the trail of our quarry. To get us started, here are a few representative samples of the idea on the moral side of things:

[Utilitarianism's] view of social cooperation is the consequence of extending to the society the principle of choice for one man, and then, to make this extension work, conflating all persons into one [...] Utilitarianism does not take seriously the distinction between persons. [Rawls (1971, p. 27).]

There are only individual people, different individual people, with their own individual lives. Using one of these people for the benefit of others, uses him and benefits the others. Nothing more. What happens is that something is done to him for the sake of others. Talk of an overall social good covers this

---

<sup>12</sup> For a few cases, see Singer (forthcoming), Ahlstrom-Vij & Dunn (2014), Goldman (2015), Dunn & Ahlstrom-Vij (2017); and for a kind of defense against empirical results reported in James (2017), see Roberts, et. al (2017).

<sup>13</sup> Rather than going this route, you might try to offer a *neutral* characterization of the epistemic separateness of propositions in terms of their alethic or probabilistic separateness (i.e., either that the propositions don't stand in entailment relations or that they are probabilistically independent). But this will not do, for obvious reasons: even if two propositions stand in the relevant entailment or probabilistic relations, haters will presumably think trading off one's truth against another is forbidden.

up. (Intentionally?) To use a person in this way does not sufficiently respect and take account of the fact that he is a separate person, that his is the only life he has. [Nozick (1974) pp. 32-33]

[Consequentialism] treats the desires, needs, satisfactions, and dissatisfactions of distinct persons as if they were the desires, etc., of a mass person. [Nagel (1970, p. 134)]

[...] [S]uppose that mankind is a super-person, whose greatest satisfaction is the objective of moral action. . . . But this is absurd. Individuals have wants, not mankind; individuals seek satisfaction, not mankind. A person's satisfaction is not part of any greater satisfaction. (Gauthier (1962), p. 126; cited in Norcross (2009).)

The basic problem with standard versions of core consequentialism is that they fail to recognize adequately the *normative separateness of persons*. Psychological autonomous beings (as well, perhaps, as other beings with moral standing) are not merely means for the promotion of value. They must be respected and honored, and this means that at least sometimes certain things may not be done to them, even though this promotes value overall. [Vallentyne (2006) p. 29; cited in Norcross (2009).]

Importantly, despite some of the misleading language (especially in the Nagel, Nozick, and Gauthier quotes), the point here is not a metaphysical one about the fact that persons enjoy distinct existences.<sup>14</sup> And, correspondingly, the fundamental mistake those who ignore the separateness of persons make is not a metaphysical one. For even supposing there were, metaphysically speaking, such a thing as a “mass-” or “super-person” comprising what we ordinarily think of as individual persons, those concerned with the moral separateness of persons would think this was neither here nor there when it came to how, morally speaking, we should treat the component persons that make up such a super-person. The mistake, in other words, is not in thinking that super-persons *exist*, but is instead in thinking that super-persons are, when it comes to what we are morally obliged to do, the appropriate *foci* of our normative thought.

Vallentyne's quote is perhaps the clearest on this point, though each of these authors makes it clear elsewhere that it is the *normative* -- not metaphysical -- separateness of persons that concerns them. As Vallentyne characterizes it, the normative separateness of persons consists in the fact that they must be “respected and honored” and not merely treated as “means for the promotion of value.” And in his later work, Rawls says that it is distinctive of persons that they are “self-authenticating sources of valid claims,” requiring our respect.<sup>15</sup> Below I'll discuss this in some more detail. For now, we can simply note that each of these authors, in different ways, understands the *normative separateness of persons* as placing restrictions on the appropriateness of certain trade-offs between what is good for one person and what is good for another. Since it's something like this latter claim that haters

---

<sup>14</sup> c.f. Norcross (2009).

<sup>15</sup> Rawls (2005, pp. 32, 72, 100)

want to make regarding the *epistemic separateness of propositions* (i.e., that there are restrictions on the appropriateness of trade-offs between them), it looks like we've found our analogue.

Let's dig in to the analogy a bit further. First notice that, in the moral domain, it is *persons* that are separate in the sense that there are restrictions on trading off what's good for one person against what's good for another. This suggests that the correct epistemic analogy is in terms of the separateness of *beliefs* not *propositions*. To explain: recall our two theories of value. Hedonism claims *persons' pleasure* has positive final practical value.<sup>16</sup> Veritism claims *true belief* has positive final epistemic value. Restrictions on trade-offs between *persons* on the basis of their separateness are restrictions on trading off one person's thing hedonism says is valuable against another person's thing hedonism says is valuable. So, if we mean to make the analogous claim in epistemology, we should not say that the issue is over whether there is a ban on trade-offs between *propositions*. Veritism doesn't say anything about the value of *propositions*. If there are restrictions on trade-offs, they are restrictions on trading off what's good for a particular belief against what's good for a separate belief. Our question about the epistemic separateness of *propositions* is, properly, a question about the epistemic separateness of *beliefs*.<sup>17</sup>

With this in mind, the idea is that, just as *persons* are separate in the sense that there are restrictions on their good being be traded off against *another person's good*, so *beliefs* are separate in the sense that there are restrictions on their *truth* being traded off against *another belief's truth*. This seems to be exactly what haters have in mind. Recall Berker's example of the atheistic scientist. What was supposed to be objectionable, according to Berker, was that in those circumstances he was considering trading off the negative value of his false belief in God against a range of other true beliefs concerning the subject matter of his scientific research.

All this should be uncontroversial. It's simply making explicit the analogue between the moral separateness of persons and the epistemic separateness of beliefs (propositions) alluded to by haters such as Berker. The interesting thing, as I'll now explain, is what happens when we take this analogue seriously.

#### **4. Interpersonal and Interdoxastic Aggregation of Value**

If persons are normatively separate in the way just described, what follows with respect to the permissibility of *trade-offs* between persons, i.e., interpersonal trade-offs? An insistence on the absolute non-negotiability of the normative separateness of persons would entail it was *never* morally permissible to sacrifice (any amount of) one person's good in order to

---

<sup>16</sup> Hedonism doesn't *quite* say this: it says that pleasure (*whoever's*) is finally practically valuable. For present purposes, we can ignore this complication and assume that the only relevant bearers of pleasure are persons. This is because, when it comes to the epistemic analog of hedonism, i.e., veritism, the only relevant bearers of *truth* are *beliefs*.

<sup>17</sup> You might think that in the present context we distinguish between token beliefs on the basis of their propositional content. Thus the "separateness of propositions" and the "separateness of beliefs" comes to one. But that will not do. For veritism says that *your* true belief that *P* is as much a bearer of epistemic value as *my* true belief that *P*.

bring about (any amount of) any number of persons' goods. In other words, insistence on the absolute non-negotiability of the normative separateness of persons entails an absolute ban on interpersonal trade-offs. But an absolute ban on the interpersonal trade-offs leads to absurd first-order results. It entails that it is *never* morally permissible to sacrifice one person's life in order to save five lives, or five million, or even five billion; a strict ban on interpersonal trade-offs entails it's never morally permissible to give one person a *short mild headache* in order to save five billion lives. Assuming, as Alistair Norcross puts it, that we don't want to follow those who insist on such a ban to the "funny farm," we should reject an *absolute* ban on interpersonal trade-offs and accept, instead, that it is at least *sometimes* morally permissible to trade-off one person's good against another's.<sup>18</sup>

Why accept the moral permissibility of interpersonal trade-offs? One reason is, as I've just said, is that rejecting it leads to absurd results -- we end up at the funny farm. Another reason, which matters more in the present context, is that *even the most ardent proponents of the moral separateness of persons accept interpersonal trade-offs*. This is worth pausing over for a moment. Rawls is often credited with his celebration of the separateness of persons. But even Rawls doesn't end up endorsing a view according to which interpersonal trade-offs are completely off the table. Even he thinks interpersonal trade-offs are *negotiable*.

Perhaps the most obvious illustration of the negotiability of such trade-offs in Rawls's theory is his commitment to a principle of distributive justice requiring that inequalities in wealth and income be justified by their positive contribution to the position of the least well-off in a society. This is his famous "Difference Principle," which says that differences in inequality in income and wealth arising from the arrangement of the basic structure of society are justified only if those differences make the least well-off in society better off than they otherwise would be under some alternative arrangement.<sup>19</sup> Implicit in the Difference Principle, then, is the permissibility of some interpersonal trade-offs. This is because the Difference Principle forbids (mere) Pareto improvements in a distribution where some people, viz. the best-off in a society, could be made better off without worsening -- but also without *improving* -- the position of the worst-off in that society. Hence the Difference Principle says it is licit -- indeed, that we are *required* -- to trade-off the good of some persons (those comprising the best-off ) against the good of some other persons (those who comprising the worse-off).

I myself don't have a theoretical stake in the moral separateness of persons. But if those who care most for and have thought hardest about the moral separateness of persons think that separateness is compatible with *at least some* interpersonal trade-offs, I'm inclined to believe them. So, here I won't be arguing any further for this idea. Let emphasize how minimal the assumption I'm making here is. Again, I'm not arguing, here, in favor of any particular conception of how, precisely, we ought to trade-off value between persons. The Difference Principle's particular claim about how to *value* trade-offs is controversial. That's not surprising, since it amounts to the claim that we should value the difference between

---

<sup>18</sup> Norcross (2009, p. 80). In fact Norcross argues for trade-offs without bound. But here I need only the weaker assumption.

<sup>19</sup> Rawls (1971, especially the discussion beginning on p. 76).



the good of the best-off after a (mere) Pareto improvement and their good prior to that improvement at *zero*. But we don't need to accept the Difference Principle's idiosyncratic claim about how to value trade-offs between persons to think it's obvious that, sometimes at least, trade-offs between persons are morally permissible and, most importantly, are not ruled out by the "separateness of persons." And I'm certainly not arguing that *unrestricted interpersonal trade-offs*, i.e., the view according to which there are no theoretical restrictions on how we can trade-off value between persons, is compatible with the separateness of persons. That would be a quite controversial assumption.

What about the corresponding claims in the epistemic case? What does the epistemic separateness of beliefs, understood analogously to the moral separateness of persons, give us reason to think about the permissibility (or not) of trade-offs? Here we come to the crux of things. For notice that, as we've just seen, a commitment to the moral separateness of persons is broadly compatible with interpersonal trade-offs. But then, by analogous reasoning, it seems we should accept trade-offs between beliefs, i.e., *interdoxastic trade-offs*.

Here is why. The epistemic separateness of beliefs is best understood as a normative matter just as the moral separateness of persons is best understood as a normative (not metaphysical) matter. In the moral case, separateness entails a set of restrictions on how persons can be treated: restrictions, in other words, on trading-off one's good against another's. But, again, *those restrictions do not constitute absolute bans* on interpersonal trade-offs. That way lies the funny farm, and in any case the proponents of the moral separateness of persons don't themselves sign up for such absolute bans on trade-offs. So, in the epistemic case, the separateness of beliefs should be understood, I'm now busy suggesting, as entailing a set of restrictions on how beliefs can be treated: restrictions, in other words, on trading-off one's truth against another's. But the idea that there are *some* restrictions on interdoxastic trade-offs does not amount to the idea that there is an *absolute ban* on them. Moreover, insofar as we take seriously the analogue from the moral case, we've good reason to think even the most ardent proponents of the epistemic separateness of beliefs will, like their counterparts in moral theory, allow for (restricted) trade-offs.

At the very least, and here's the important point, we've been given no reason by haters to think the situation is substantively different in the epistemic case than it is in the moral. Worse, it appears that the best way of cashing out the otherwise obscure metaphor regarding the separateness of beliefs, viz. by analogy with the moral separateness of persons, gives us positive reason for thinking trade-offs will sometimes be permissible. In fact, since haters *themselves* encourage the analogy between the two domains, they thereby encourage the thought that what's goes for trade-offs between morally separate persons (they're permissible, with restrictions) goes for trade-offs between epistemically separate beliefs (they're permissible, with restrictions). Haters therefore face a dilemma.

On the one hand, they can have a plausible conception of the epistemic separateness of beliefs. The most plausible conception -- suggested by haters themselves, and modeled on the moral separateness of persons -- gives us no independent reason for thinking that there is an *absolute ban* on trade-offs between beliefs. If we understand the epistemic

separateness of beliefs as the epistemic analogue of the moral separateness of persons, we should accept that trade-offs will sometimes be epistemically permissible.

On the other hand, haters can reject the conception of the epistemic separateness of beliefs modeled on the moral separateness of persons. I haven't argued, here, that this rejection would be unworkable or even unreasonable. If haters reject the conception of the epistemic separateness of beliefs modeled on the moral separateness of persons, they cannot then remain silent on what the epistemic separateness of beliefs is supposed to amount to. Such metaphors are meaningless unless unpacked. And, importantly, it cannot turn out that what it is for beliefs to be epistemically separate is for it to be such that their truth cannot ever be appropriately traded-off against another belief's truth. That would beg the question against the proponent of the package of EOC and veritism. I'm not optimistic about the prospects for haters giving a meaningful alternative conception of beliefs' epistemic separateness that doesn't simply beg the question in this way; but whether it can be done remains to be seen: if they take this route, the burden is now on haters to tell us what they mean.

This might seem unfair to haters. Berker, at least, is explicit about the fact that he thinks there are structural disanalogies between the moral and the epistemic domains.<sup>20</sup> And it's reasonable to assume Berker thinks it is these disanalogies that explain why the moral separateness of persons is negotiable and therefore allows for trade-offs whereas the epistemic separateness of beliefs is not, and so does not. After all, it's Berker himself who points to the disanalogy in setting up his claim initially. As he says, *unlike in the moral case*, "in the epistemic case there is no wiggle room."<sup>21</sup> But what I am presently urging is that we are owed an explanation for this lack of wiggle room. And the appeal to the metaphor of the separateness of beliefs (or propositions) is no help at all. At best, it confuses matters.

## 5. Cases and Cases

Still, you might have worries. Perhaps I'm placing far too much emphasis on a catchy turn of phrase. (In my defense, I'm trying here not just to bury the phrase but to make sure it's dead when I do.) This can seem unfair to the hater. What she cares about isn't defending her use of a bit of language (the "separateness of propositions"), but instead her view about *trade-offs* (they're always forbidden). And after all, the worry continues, Berker has an explanation for the lack of wiggle room in the epistemic case that doesn't appeal to anything about separateness: according to him, epistemic rationality is *non-promotive*, in the sense that it's not concerned with promoting some valuable end such as true belief. So, what's important is not the epistemic separateness of beliefs (propositions), whatever that turns out to mean. What's important (a hater might say) is why we should think that epistemic rationality is non-promotive. So (says our hater), forget about arguments over separateness; instead, focus on arguments over whether epistemic rationality is promotive. In this section, I address this issue.

So: Why do haters think that epistemic rationality is non-promotive, i.e., not concerned with promoting some independently valuable end such as true belief? The answer is precisely

---

<sup>20</sup> Berker (2013a, 2013b, 2015).

<sup>21</sup> Berker (2013a, p. 365).

the one Berker gives: because of our judgments about cases. Recall, again, haters' intuitions, which we're invited to share, about cases such as Berker's:

**Scientist:** Suppose I am a scientist seeking to get a grant from a religious organization. Suppose, also, that I am an atheist: I have thought long and hard about whether God exists and have eventually come to the conclusion that He does not. However, I realize that my only chance of receiving funding from the organization is to believe in the existence of God: they only give grants to believers, and I know I am such a bad liar that I won't be able to convince the organization's review board that I believe God exists unless I genuinely do. Finally, I know that, were I to receive the grant, I would use it to further my research, which would allow me to form a large number of new true beliefs and to revise a large number of previously held false beliefs about a variety of matters of great intellectual significance.

According to Berker, "most of us" would not think that a belief in God was reasonable, or rational, or justified, in these circumstances. As we know, haters' point isn't about how to *value* trade-offs: their point is not that in this case, all things considered, it would be unreasonable to believe in God in order to further one's research (because, say, the research furthered isn't sufficiently important). Instead, their point is that such trade-offs are *always forbidden*. It's important to be clear on this point, so I'll pause over it for a moment.

It can help to put things in terms of reasons. Haters' point is not that the epistemic reasons in favor of believing in God (because it will further one's research) are *outweighed* by the epistemic reasons against doing so (because, say, it's something one thinks is false). Instead, haters' view is that the epistemic reasons in favor of believing in God in this case *do not exist*. We know this is haters' view because their central thought isn't one about how to *value* trade-offs -- about how much (important) truth we'd have to put up against how much (important) falsehood for things to come out equals-pequals. If that were haters' central thought, then it'd be utterly mysterious why they think such trade-offs are always forbidden. After all, it'd be possible to imagine the sort of case where the scales tip in favor of the trade-off. But, and again, haters think such trade-offs are always forbidden. In the language of the misleading metaphor from the previous section: the epistemic separateness of beliefs is *non-negotiable*.

But as I'll now explain, this way of arguing for the non-promotive nature of epistemic rationality, and so in favor of an absolute ban on trade-offs, via our judgments about cases, suffers a serious methodological mistake that vitiates its probative value. The mistake is one Mark Schroeder identifies in a different context, and it'll be helpful to work our way into the mistake as made by haters by first looking at the instance identified by Schroeder:<sup>22</sup>

---

<sup>22</sup> In unpublished work, Daniel Wodak marshals this same diagnosis of a symmetrical methodological mistake in a different domain, viz. arguments regarding the reasons we have on the basis of normative (and in particular aesthetic) testimony. I am in complete agreement with what he says there, which is, roughly, that arguments for the conclusion that we lack reasons for attitudes on the basis of normative testimony that go by way of our intuitive verdicts about the wrong (or no) comparison cases are not probative.

In a first version of the case, you see Tom Grabit come out of the library, pull a book from beneath his shirt, cackle gleefully, and scurry off. Intuitively, you have a reason to believe that Tom just stole a book from the library. But in a revised version of the case, Tom has an identical twin, Tim, from whom you cannot visually distinguish him. If you're aware of this, then it turns out that you don't have any reason to believe that Tom stole a book after all. Right? Of course that's right. This, after all, is a classic case of an undercutting defeater, and as everyone knows, undercutting defeaters make it the case that things that would otherwise have been reasons for you, instead are not.

This is a natural thought about the case, but your intuitions mislead you. That you still do have some reason to believe that Tom stole a book can be observed by comparison with yet a third version of the case. In the third version, Tom and Tim have a third identical sibling, Tam. In this case, you have even less reason to believe that Tom stole a book than in the second, and so in the second it can't have gone away entirely. By similar reasoning, you still have a reason to believe that Tom stole a book even in the third case, because there is a fourth case in which there are four identical siblings and your reason to believe that Tom stole a book is still worse.<sup>23</sup>

Schroeder's point explicitly concerns undercutting defeat. But in the present context, what's important is the more general lesson we can draw from Schroeder's examples, viz. that we can be led into intuitive error about the (non-) existence of (epistemic) reasons by failing to draw the right comparisons between cases -- either by using the *wrong* comparison case or by failing to use *any comparison case whatsoever*.<sup>24</sup> That, in short, is the mistake haters make. Let me explain.

If the lesson of Schroeder's Grabit example is that our intuitive verdicts regarding the existence of epistemic reasons can be misled by failing to use the correct comparison case, then failing to use *any comparison case whatsoever* in coming to an intuitive verdict regarding the existence of epistemic reasons -- as we're asked to do in, e.g., Berker's atheistic scientist case, might equally well mislead. And that, I think, is exactly what's happened. To see this, consider the following case, which I'll call:

**Inept Scientist:** Suppose I am a scientist seeking to get a grant from a religious organization. Suppose, also, that I am an atheist: I have thought long and hard about whether God exists and have eventually come to the conclusion that He does not. However, I realize that my only chance of receiving funding from the organization is to believe in the existence of God: they only give grants to believers, and I know I am such a bad liar that I won't be able to convince the organization's review board that I believe God exists unless I genuinely do. Finally, I know that, were I to receive the grant, I would

---

<sup>23</sup> Schroeder (2004, p. 93).

<sup>24</sup> Schroeder himself uses the case to draw a lesson about our intuitions regarding reasons quite generally, including our intuitions about *practical* reasons. See Schroeder (2004, esp. pp. 93 and following) for discussion.

use it to further my research, which would allow me to form a small number of new true beliefs and to revise a small number of previously held false beliefs about a variety of matters of minimal intellectual significance.

Note that Inept Scientist is exactly like Berker's Scientist, except that, whereas in Scientist the grant from the religious organization allowed me to form a *large* number of new true beliefs and revise a *large* number of previously held false beliefs about a variety of matters of *great* intellectual significance, in Inept Scientist, I form a *small* number of new true beliefs and revise a *small* number of false beliefs about matters of *minimal* intellectual significance. I hereby invite you to think (presumably in agreement with haters) that forming a belief in God's existence would not be epistemically reasonable in these (Inept Scientist) circumstances. But that all-out verdict, that it would be epistemically unreasonable to form a belief in God's existence in these circumstances, isn't the verdict I'm interested in exploring.

Instead, consider the essentially *comparative* judgment about Inept Scientist as opposed to Scientist. We should agree that it is *less* reasonable to believe in God's existence in Inept Scientist than it is in Scientist. If you're not convinced, consider a *third* version of the case that is exactly like Inept Scientist except that one doesn't form *any* new true beliefs (and so only revises a small number of previously held false beliefs about matters of minimal intellectual significance). Or consider a *fourth* version of the case that is just like the third version but where still *fewer* false beliefs are revised. In each of these successive cases, the correct intuitive verdict seems to be that one is progressively *less justified* or *less reasonable* in forming a belief in God's existence in the circumstances one finds oneself in.

Let's put this again in the language of reasons. The intuitive verdict about Inept Scientist is that you have *less* reason to believe in God's existence in those circumstances than you do in Scientist. And similarly for the third and fourth versions of the case: these are cases where you have progressively less and less reason to believe in God's existence. But if this is so, then in Scientist you must have *some* reason to believe in God's existence. This is because in Inept Scientist we know that your reason to believe hasn't gone away entirely; and similarly for the third and fourth versions of the case: you still have reason to believe in God's existence in Inept Scientist because in the third (and fourth) version of the case you have *less* (and still less). And so examples such as Scientist do not put the lie to a promotive conception of epistemic rationality, since when we consider such examples in comparison to other cases, we see that there *is* some epistemic reason in such cases, which is what the promotive conception of epistemic rationality predicts.

In sum, then: haters' view that epistemic trade-offs are always forbidden is not meant to be tied to verdicts about all-things-considered rationality in particular cases. Their view is that there is an *absolute* ban on trade-offs because epistemic rationality is non-promotive. So, it's not supposed to be essential to the force of the examples meant to support this view that we agree with haters about the all-things-considered verdict in any given case. What's supposed to be essential is the intuitive verdict that in the relevant cases there is no epistemic reason *at all* in favor of the belief in virtue of (potential) trade offs. Let's suppose we grant, for the sake of argument, that we share this intuitive verdict in one-off cases. This

does not show, what it is supposed to show, that we share haters' view that epistemic rationality is non-promotive. This is because when we *contrast* our verdict about such cases with suitably modified versions of the cases, the intuitive verdict there is that there is *still less* reason in these modified cases -- even if the *overall* verdict (i.e., the judgment about whether the trade-off is here permissible, whether it would be reasonable or justified) is the same. But if there is still less reason in the modified cases, then there must be *some* reason in the unmodified cases, contrary to our initial intuition. That is Schroeder's lesson. So, in the cases (modified and unmodified alike) there must be some reason, in virtue of the fact that doing so would promote true belief (and promote the avoidance of false belief), to believe. And so epistemic rationality is (at least in part) promotive, contrary to haters' claim.

I'll close by considering one final rejoinder. Perhaps this is still unfair to haters -- and especially to Berker -- since it might seem to rely on the details of particular examples, and in other work Berker explicitly eschews such reliance as the right way to understand his arguments:

I do not mean my argument to rest too heavily on these particular examples. The important point is not the examples themselves, but rather the structure of the examples. As I see it, my [...] argument against epistemic consequentialism involves not a series of counterexamples, but rather a *recipe* for generating counterexamples. [...] this recipe makes my style of argument particularly versatile: even if fault is found with the specific examples I have offered in this essay, my recipe allows me to find other examples of a structurally similar sort that suffice to make the same point.<sup>25</sup>

But what I've just argued is that there is a *structural flaw* in Berker's recipe, and so it can't be used to cook up the results he wants. He's failed to recognize that some ingredients in it, to taste right in the cake, are essentially comparative: the sugar must be weighed *against* the flour. Less culinarily: our intuitive verdict about any single case that believing *P* is unreasonable despite the fact that so believing would best promote overall epistemic value doesn't show that epistemic rationality is non-promotive. This is because, as we just saw, when we compare our verdict about such a case to one where believing *P* would *less well* promote the epistemic values, our judgment is that it would be *even more unreasonable* to believe *P* in that case. And that means, again as we just saw, that it must be that there is *some* reason to believe *P* in the first case -- and so epistemic rationality must have *some* promotive component (for that is what explains the existence of the reason). At best, then, what our intuitive verdicts about these singleton, non-comparative, cases suggest is that there are *some* restrictions on epistemic trade-offs; our intuitive verdicts don't deliver the result that such trade-offs are absolutely forbidden.

## 6. Summary

Haters object to the pairing of veritism and EOC on the grounds that it ignores the epistemic separateness of beliefs (propositions) and thereby allows for objectionable

---

<sup>25</sup> Berker (2013b p. 378; emphasis in original).

epistemic trade-offs. I've argued here the metaphor of separateness of beliefs is at best inert, at worst misleading. It's misleading if unpacked by analogy with the moral separateness of persons. This is because, in that domain, even the most ardent proponents of the separateness of persons don't think that separateness yields a ban on trade-offs. It's inert if it remains a metaphor: it can't do any argumentative work.

If the argumentative work is done elsewhere, it appears to be done by appeal to our intuitive verdicts about cases such as Berker's Scientist. But those verdicts are not probative, for they fail to recognize that our intuitive verdicts about the existence of epistemic reasons in particular cases can mislead when those verdicts aren't tested against the right (or indeed *any*) comparison cases. Once we introduce the relevant comparison cases, we see that our intuitive verdicts tell in favor of a promotive conception of epistemic rationality, i.e., one where trade-offs are at least in principle permissible.

Let me close with a remark that's more confession than argument, but which is short. It makes pre-theoretical sense to me to think the activity of trading off persons' goods requires special justification. After all, we're dealing with a *person* and her welfare; of *course* special justificatory caution is required in circumstances where we're deliberately considering *harming* a *person* (even when it's to the benefit of others). At least, I have a sense for why I might feel this way, pre-theoretically, even if I end up in a position where I say otherwise. Why might this be? For one thing, and quite obviously, *I* am a person, and *I* do not particularly want to suffer harm even for others' benefit. For another, many of those who I care about most are *also* persons, and I do not much want that *they* be cut up to save the many. And even if I end up where the consequentialist in the practical domain wants me to be, it's important that I recognize the need to be *talked into it*. But, and maybe this is just my cold and unsympathizing epistemic nature, I'm at a loss to make any good pre-theoretical sense out of the idea that trading off one belief's truth against another's requires, in a similar way, any sort of special justification. After all (I'm inclined to say) we're (just) dealing with a *belief*. Sure, the thought continues, I'm considering *believing falsely* in order to believe truly about some other matters. But it is just belief, just some mental state of mine; *I* am not the belief, nor are any who I care about most themselves beliefs. So, unlike in the practical case, I don't feel the need to be *talked into* the position the epistemic consequentialist wants me to be in: I can happily start from there, and need to be convinced otherwise.

## References

- Ahlstrom-Vij, K. (2013), In Defense of Veritistic Value Monism. *Pacific Philosophical Quarterly*, 94: 19–40.
- Ahlstrom-Vij, K. & Dunn, J. (2014). A Defence of Epistemic Consequentialism. *Philosophical Quarterly* 64 (257):541-551.

- Andow, James (2017). Do non-philosophers think epistemic consequentialism is counterintuitive? *Synthese* 194 (7):2631-2643.
- Berker, Selim (2013a). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review* 122 (3):337-393.
- Berker, Selim (2013b). The Rejection of Epistemic Consequentialism. *Philosophical Issues* 23 (1):363-387.
- Berker, Selim (2015). Reply to Goldman: Cutting Up the One to Save the Five in Epistemology. *Episteme* 12 (2):145-153.
- Bradley, Ben (2006). Two Concepts of Intrinsic Value. *Ethical Theory and Moral Practice* 9 (2):111-130.
- Bradley, Ben (2009). *Well-being and Death*. Oxford University Press.
- Dunn, Jeffrey & Ahlstrom-Vij, Kristoffer (2017). Is Reliabilism a Form of Consequentialism? *American Philosophical Quarterly*.
- Feldman, Fred (2004). *Pleasure and the Good Life: Concerning the Nature, Varieties and Plausibility of Hedonism*. Clarendon Press.
- Firth, Roderick (1978). "The Schneck Lectures, Lecture 1: Epistemic Utility." In Firth 1998.
- Firth, Roderick (1998). *In Defense of Radical Empiricism: Essays and Lectures*. ed. John Troyer. Lanham, MD: Rowman and Littlefield.
- Fumerton, Richard (2001). "Epistemic Justification and Normativity." In Steup 2001, pp. 49 - 60.
- Gauthier, David (1962). *Practical Reasoning: The Structure and Foundations of Prudential and Moral Arguments and Their Exemplification in Discourse*. Oxford: Oxford University Press.
- Goldman, Alvin I. (1999). *Knowledge in a Social World*. Oxford University Press.
- Goldman, Alvin I. (2015). Reliabilism, veritism, and epistemic consequentialism. *Episteme* 12 (2):131-143.
- Heathwood, Chris (2006). Desire Satisfactionism and Hedonism. *Philosophical Studies*, 128(3), 539-563



- Horowitz, Sophie (forthcoming). Epistemic Value and the Jamesian Goals. In Jeffrey Dunn Kristoffer Ahlstrom-Vij (ed.), *Epistemic Consequentialism*. Oxford University Press.
- Jackson, Frank & Pargetter, Robert (1986). Oughts, options, and actualism. *Philosophical Review* 95 (2):233-255.
- Korsgaard, Christine M. (1983). Two distinctions in goodness. *Philosophical Review* 92 (2):169-195.
- Littlejohn, Clayton (forthcoming). The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology. In Kristoffer Ahlstrom-Vij Jeff Dunn (ed.), *Epistemic Consequentialism*. Oxford University Press.
- Littlejohn, Clayton (2015). Who Cares What You Accurately Believe? *Philosophical Perspectives* 29 (1):217-248.
- Mill, J.S. (1861). *Utilitarianism*. Hackett Publishing Inc.
- Nagel, Thomas (1970). *The Possibility of Altruism*. Oxford Clarendon Press.
- Norcross, Alastair (2009). Two dogmas of deontology: Aggregation, rights, and the separateness of persons. *Social Philosophy and Policy* 26 (1):76-95.
- Nozick, Robert (1974). *Anarchy, State and Utopia*. Basic Books.
- Pritchard, D. (2010). 'Knowledge and Understanding', in *The Nature and Value of Knowledge*, eds. D. Pritchard, A. Millar, A. Haddock, 3-88. Oxford: Oxford University Press.
- Rawls, John (1995). *Political Liberalism*. Columbia University Press.
- Rawls, John. (1971). *A Theory of Justice*. Belknap.
- Roberts, Pendaran ; Andow, James & Schmidtke, Kelly (2017). Lay intuitions about epistemic normativity. *Synthese*:1-21.
- Schroeder, Mark (2004). *Slaves of the Passions*. Oxford: Oxford University Press.
- Singer, Daniel J. (forthcoming). How to be an Epistemic Consequentialist. *Philosophical Quarterly*.
- Steup, Matthias, ed. (2001). *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press.
- Vallentyne, Peter (2006). Against Maximizing Act Consequentialism. in James Dreier, ed., *Contemporary Debates in Moral Theory* (Oxford: Blackwell, 2006).

Zagzebski, L. (2004). 'Epistemic Value Monism', in Ernest Sosa: and his Critics, ed. J. Greco, 190-198. Oxford: Wiley-Blackwell.

Zimmerman, Michael J. (2001). *The Nature of Intrinsic Value*. Rowman & Littlefield.