The Journal of Ethics (2020) 24:177–192
https://doi.org/10.1007/s10892-019-09317-y

Epicureanism and the Wrongness of Killing

Tim Burkhardt[1]

Received: 19 September 2019 / Accepted: 16 December 2019 / Published online: 27 December 2019

Abstract
This paper argues that Epicureanism about death is consistent with grounding the wrongness of killing in the interests of the victim. Both defenders and critics of Epicureanism should agree that, if we knew Epicureanism to be false, then we would have a moral reason not to kill people. We would have this reason because we would know that killing people harms them. And even Epicureans should agree that, given their evidence, Epicureanism could be false. Given that it could be false, and given that we would harm people by killing them if it were, we in fact have a moral reason not to kill them—a reason which, as this paper will show, is both grounded in their interests and consistent with the failure of death to be in any way bad. The latter part of the paper discusses some advantages that this approach enjoys over two other attempts to reconcile Epicureanism with the wrongness of killing, by David Hershenov and Mikel Burley, respectively.

Keywords Badness of death · David Hershenov · Epicureanism · Killing · Mikel Burley · Moral reasons

1 Introduction

Epicureanism about death is widely thought to threaten our conviction that killing people is *prima facie* seriously wrong.[1] The worry is a natural one: If death is not bad for the one who dies, then the obvious grounds for condemning killings are no longer available (Hershenov 2007: 171). At best, we might explain the wrongness of killing by appealing to other-affecting considerations—effects on survivors, on overall utility, or its tendency to brutalize the killers, for example. But these sorts of explanations have severe limitations, and even when they do deliver the right

[1] For example, see Feit (2016: 143), Pleasants (2008: 261), Silverstein (1980: 413) and Yourgrau (2000: 55).

✉ Tim Burkhardt
tim.burkhardt@duke.edu

[1] Department of Philosophy, Duke University, 201 West Duke Building, Box 90743, Durham, NC 27708, USA

verdict, they do so for reasons that are seriously incomplete (Burley 2010: 72; Marquis 1989: 189). The effects they appeal to are certainly significant, but they do not seem to be the central reason for condemning killing. That reason seemingly has to be one that is grounded in the interests of the *victim.* (Marquis 1989: 189; McMahan 2002: 95). The challenge for Epicureans is to account for this in a way that is consistent with their view.

This paper attempts to meet that challenge. Both defenders and critics of Epicureanism should agree that, if we knew Epicureanism to be false, then we would have a moral reason not to kill people. We would have this reason in virtue of the harm that killing would inflict on them (I assume throughout that all lives are worth living). And even dyed-in-the-wool Epicureans ought to concede that, given our evidence, Epicureanism could be false. Given that it could be false, and given that we would have a moral reason not to kill people if it were, we in fact have such a reason.[2] The goal of what follows is to set out this argument in detail, defend its premises, and then consider some objections to it. The latter half of the paper will then contrast it with some other attempts to reconcile Epicureanism with the wrongness of killing.

2 Preliminaries

First some useful distinctions.

My argument makes claims about our practical reasons. These are reasons to engage in or refrain from certain kinds of behavior: my reasons to write this paper, to wash my dishes, or to refrain from smoking, for example. Practical reasons contrast with epistemic reasons—reasons for or against believing something—such as my reason to doubt that this paper will be widely read, and causal reasons, such as the reasons for last week's bad weather.

The goal of my argument is to show that Epicureanism is consistent with a particular kind of practical reason to refrain from killing: a *moral* reason. Following Michael Huemer, I take moral reasons to be distinguished from other practical reasons by two main characteristics. The first is that they are non-selfish. Their force does not depend on whether the behavior they favor promotes the agent's own interests (Huemer 2013: 260). There may be exceptions to this. For example, I can imagine a moral reason to refrain from self-harm which is grounded in the interests of the agent (who, in this case, is also the victim). But even if there are such cases, moral reasons can still be distinguished by the fact that they at least *can* be non-selfish. That sets them apart from prudential reasons, for example, which must always derive their force from the agent's own interests.

The second distinguishing characteristic of moral reasons is that they are categorical. Their force does not depend on whether the behavior they favor satisfies the agent's desires (Huemer 2013: 261). Again, we might be able to imagine exceptions. Perhaps an agent could have a moral reason to satisfy his own desires if doing so

[2] This argument is inspired by Michael Huemer's (2013) proof of moral realism.

would lead to higher total utility, or something like that. But even if this is possible, it can still be true that moral reasons do not need to derive their force from the agent's desires.

Like practical reasons in general, moral reasons can be either *pro tanto* or all-things-considered reasons. If some factor counts in favor of my doing something, then it gives me a *pro tanto* reason to do it; if it counts against my doing something, it gives me a *pro tanto* reason to refrain from doing it. Whether or not I have an all-things-considered reason to behave in some way depends on how my *pro tanto* reasons for and against that behavior stack up against each other. For example, impressing my girlfriend gives me a *pro tanto* reason to use steroids, but the associated health risks give me a *pro tanto* reason not to. If I am right to judge that the risks of impressing her outweigh the benefits,[3] and if there are no other reasons in play, then I have an all-things-considered reason to refrain from using steroids. This example involves prudential reasons, but the same idea applies to moral ones. For example, if flipping a switch would save the lives of five track workers, then that gives me a *pro tanto* moral reason to flip it; that it would kill a sixth worker gives me a *pro tanto* moral reason not to. If the former reason outweighs the latter, and if these are the only reasons in play, then I have an all-things-considered moral reason to flip the switch. In this paper, I argue that Epicureanism is consistent with a *pro tanto* moral reason to refrain from killing.

The next distinction I want to make is between objective and subjective practical reasons. As Huemer explains, objective practical reasons are those which "obtain in virtue of the actual circumstances surrounding the agent, regardless of whether the agent knows of them or has any reason to believe that they obtain" (Huemer 2013: 260). In contrast, subjective reasons "determine what it is rational to do, or what it makes sense to do, from the agent's perspective, or, given what the agent is aware of at the time of the decision making." He illustrates the difference with the case of Thurston, a thirsty man with a justified belief that the glass in front of him contains drinkable water. In fact, it contains an odorless form of petrol. Given his thirst, Thurston has a subjective reason to drink from the glass, but no objective reason to do so (Huemer 2013: 259). Here is another example, this one involving moral reasons: A hungry child asks for a bite of your peanut butter sandwich. You justifiably but wrongly believe that this child has a severe peanut allergy. You therefore have a subjective moral reason to refrain from sharing, but no objective moral reason to refrain.[4] This paper argues that Epicureanism is consistent with a *pro tanto*, subjective moral reason to refrain from killing—subjective because, as we will see, it is grounded partly in the evidence we have.

[3] She would probably not be very impressed anyway.

[4] These examples describe divergences of subjective and objective reasons, but when an agent's evidence matches the actual states of the world, then any reason grounded in that evidence is also grounded in those states, and is therefore both objective and subjective. For example, if Thurston's glass were in fact filled with potable water, as he justifiably believes it to be, then his reason to drink would be both objective and subjective. So these are not mutually exclusive categories.

Lastly, I want to distinguish between moral reasons according to the sources of their force. I call moral reasons *victim-affecting* when they derive their force from the interests of the target moral patient—the *victim*.[5] This contrasts with moral reasons deriving their force from considerations of total utility or possible effects on bystanders, for example, which we can call *utility-* and *other-affecting* moral reasons, respectively. Sometimes these reasons overlap. A moral reason not to abuse people is a good example. It is victim-affecting because abuse is harmful to the abused person, but also utility- and other-affecting because it lowers total utility and harms those who witness it. In other cases, victim-affecting reasons will diverge from utility- and other-affecting reasons. What we have all-things-considered reason to do in cases of divergence will depend on how the various reasons in play stack up against each other. That is a hotly debated question but not one that we need to answer here, since the reason for which I will argue is only *pro tanto*.

The distinction between victim-affecting moral reasons on the one hand and utility- and other-affecting ones on the other is necessary because of the nature of the worry about Epicureanism. The worry is not that Epicureanism is inconsistent with *every* moral reason to refrain from killing. Even anti-Epicureans should agree that it is not, for it has no problem offering utility- and other-affecting reasons not to kill. Rather, the worry is that Epicureanism is inconsistent with the most important reasons to refrain from killing, namely those grounded in the interests of the victim. This is understandable. After all, to kill is to cause death, so if death is not bad for the one who dies, then killing causes no harm to the victim. And if it causes no harm to the victim, then it is difficult to see how a moral reason to refrain from killing could be grounded in the victim's interests. This worry is not helped by appeals to utility- and other-affecting moral reasons. Instead, what needs to be shown is that Epicureanism is consistent specifically with a victim-affecting moral reason not to kill.

This is what I will now try to show.

3 The Argument, Premise by Premise

With the above distinctions on the table, let me offer a more precise statement of my argument:

1. If we knew Epicureanism to be false, then we would have a *pro tanto*, subjective, victim-affecting moral reason not to kill anyone.
2. Given our evidence, Epicureanism could be false.

[5] Depending on the kind of conduct at issue, 'victims' are more appropriately called 'beneficiaries'. For example, we have a victim-affecting moral reason to be kind to others: victim-affecting because it is a reason grounded in their dignity and well-being. 'Victim' sounds admittedly bizarre in such a context. But I will retain the term anyway, both for simplicity and because this paper deals with reasons not to kill, in which context it is certainly appropriate.

3. Given (1) and (2), we have a *pro tanto*, subjective, victim-affecting moral reason not to kill anyone.

Therefore,

4. We have a *pro tanto*, subjective, victim-affecting moral reason not to kill anyone.

This section defends each premise.

3.1 If We Knew Epicureanism to be False, Then We Would Have a Pro Tanto, Subjective, Victim-Affecting Moral Reason not to Kill Anyone

If we knew Epicureanism to be false, we would know that death is harmful to the one who dies. To kill is to cause death. So, if we knew Epicureanism to be false, we would know that killing harms its victim. That would give us a reason not to kill. And since this reason derives its force from the victim's interest in avoiding harm, that makes it a victim-affecting moral reason: we would have it even if we enjoyed killing and even if the victim's death would be in some way good for us. Moreover, this reason not to kill would be subjective, since it would be grounded in whatever evidence is responsible for our knowledge that Epicureanism is false.[6] Thus, if we knew Epicureanism to be false, then we would have a subjective, victim-affecting moral reason to refrain from killing. Perhaps this reason could be outweighed if killing were necessary for self-defense or for the saving of many innocent lives, for example. But even then, it would always count against killing to at least some degree. That makes it a *pro tanto* reason. So, if we knew Epicureanism to be false, we would have a *pro tanto*, subjective, victim-affecting moral reason not to kill anyone.

3.2 Given Our Evidence, Epicureanism Could be False

This is a claim about epistemic probability. The epistemic probability of a proposition is a function of the justification we have for thinking that it is true: the more justification we have, the greater the epistemic probability (Huemer 2013: 262). A total absence of justification yields an epistemic probability of zero, and a conclusive proof yields an epistemic probability of one. What this premise implies, therefore, is that we have no conclusive proof of Epicureanism. Given our evidence, the epistemic probability that Epicureanism is true is less than one. How much less will depend on the precise nature of our evidence, and that will vary from person to person. Epicureans will judge the probability to be somewhere between one half and one, and anti-Epicureans at less than one half. But absent a conclusive proof,

[6] This is not to say that the reason is not also objective. Indeed, since knowledge entails truth, the evidence responsible for our knowledge of Epicureanism's falsity would match the actual states of the world. So any reason grounded in that evidence would also be grounded in those states, and would therefore be both objective and subjective (see n. 4, above).

everyone ought to judge the probability to be less than one. Therefore, everyone ought to judge the probability that Epicureanism is false to be greater than zero.

Suppose that we did have a conclusive proof of Epicureanism. That would refute this premise. But it would also make my project unnecessary. If we had a conclusive proof of Epicureanism, then we could safely ignore every objection to that view, including those related to the wrongness of killing. Anyone interested in pressing such an objection will therefore grant this premise.

3.3 Given (1) and (2), We Have a Pro Tanto, Subjective, Victim-Affecting Moral Reason not to Kill Anyone

This is the key premise of the argument. It is not as straightforward as the preceding two, but it is based on a very simple and intuitive principle. The principle is this: if knowledge of some proposition, P, would give us a subjective reason to behave in some way, Φ, then, to the extent that we have reason to believe that P, we have a subjective reason to Φ.[7] Perhaps the strength of this reason to Φ depends on the strength of our reason to believe that P—that is, on the epistemic probability that our evidence allows us to assign to P. If so, then the lower the epistemic probability of P, the weaker is our reason to Φ.[8] But no matter how weak our reason to Φ is, it will always count in favor of Φ-ing to at least some degree.

Here is an example to make this less abstract. Imagine that Finnegan the football fan is watching the coin toss at the beginning of a game. The captain of one of the teams, the Roid Rage Rascals, calls "heads". Finnegan knows that, whenever the Rascals win the coin toss, they elect to receive the ball first. Thus, if Finnegan knew that the coin will land heads, then he would have a reason to believe that the Rascals will receive the ball first. Now, prior to the toss, Finnegan does not know whether the coin will land heads. But he already has at least some reason to believe that it will. It is a fair coin and fair coins land heads half of the time. So, given Finnegan's evidence, the epistemic probability that the coin will land heads is greater than zero. According to the above principle, Finnegan therefore has at least some reason to believe that the Rascals will receive the ball first. More precisely, he has a *pro tanto*, subjective reason to believe that they will receive it first: *pro tanto* because it can be outweighed (by his later knowledge of how the toss actually turns out, for example) and subjective because it is based on his evidence (that the coin is fair and that the Rascals will choose to receive first if they win the toss, for example).

Of course, this does not mean that Finnegan is justified in believing that the Rascals will receive the ball first. Since the coin is fair, they are just as likely to win the toss as they are to lose it. And if they lose it, then their rivals, the Meathead

[7] Compare Huemer's Probabilistic Reasons Principle (Huemer 2013: 263–265).

[8] This sort of covariation is plausible, but it is not strictly necessary. The principle is compatible with any view about how the strength of our reason to Φ is affected by the epistemic probability of P. For instance, it allows that the reason to Φ is equally strong whenever the epistemic probability of P is above a certain threshold. In the extreme case, that threshold could be set at zero. Our reason to Φ would then always be of equal strength as long as our evidence leaves even the slightest chance that P is true.

Monkeys, will elect to receive first. Finnegan's reason to believe that the Rascals will receive first is therefore no stronger than his reason to believe that the Monkeys will. Again, it is only a *pro tanto* reason, not an all-things-considered one. And since it is just as strong as his reason for believing that the Monkeys will receive first, I assume that Finnegan is not justified in holding either belief. But for all that, he still has at least some reason to believe that the Rascals will receive first.

Here is a second example, this one involving moral reasons. Imagine that Paulo and his friends are celebrating his birthday with a piñata. Paulo is blindfolded and given a bat, and his friends egg him on as he prepares to take a swing. Here is a proposition: If Paulo swings the bat, his friend Dolores will be injured. Call this proposition DI. If Paulo knew DI to be true, he would have a reason to refrain from swinging the bat. And Paulo does have at least some reason to believe that DI is true. He knows that Dolores is in attendance, but because he is blindfolded, he cannot tell exactly where she is. Perhaps she is somewhere within the arc that his bat would trace if he were to swing it. Paulo cannot rule this out. So, given his evidence, the epistemic probability of DI is greater than zero. Paulo therefore has a reason to refrain from swinging the bat. To be precise, he has a subjective, victim-affecting moral reason to refrain from swinging it: subjective because it depends on his evidence,[9] and victim-affecting and moral because it is grounded in Dolores' interest in avoiding harm rather than Paulo's interests or desires.

Of course, this does not mean that Paulo has an all-things-considered moral reason to refrain from swinging the bat. That depends on what other reasons are in play. Suppose, for example, that Paulo has a competing subjective reason that counts in favor of swinging. This reason is grounded in his evidence for the following proposition, H: Paulo has a desire to hit the piñata. Suppose that H is true. In that case, Paulo is likely to have at least some evidence for it, and therefore a subjective reason to swing the bat.[10] Now, on the assumption that no other subjective reasons are in play, what Paulo has an all-things-considered subjective reason to do depends on which of the these two reasons is stronger—the one grounded in his evidence for DI, or the one grounded in his evidence for H. Plausibly, the answer to that depends on two factors: (i) the strength of the evidence in which these reasons are grounded and (ii) the strengths of the corresponding objective reasons. The first factor is straightforward: The more evidence Paulo has for H, the stronger is his reason to swing; the more he has for DI, the stronger is his reason to refrain from swinging.[11]

As for the second factor: If DI and H are both true, then they each ground an objective reason which corresponds to one of the subjective reasons grounded in Paulo's evidence. In particular, if DI is true, then it grounds an objective reason to refrain from swinging which corresponds to the subjective reason grounded in Paulo's evidence for DI. And likewise for H: if true, then it grounds an objective reason

[9] Specifically, on its failure to rule out DI.

[10] Indeed, since Paulo can presumably discover the truth of H by introspection, his evidence for it, given that it is true, is likely to be very strong.

[11] Although this is plausible, it is not strictly necessary. We could reject it without harm to my argument. See n. 8.

to swing which corresponds to the subjective reason grounded in Paulo's evidence for H. Plausibly, the strengths of these subjective reasons depend not only on the strength of Paulo's evidence for DI and H, but also on the strengths of the corresponding objective reasons: the stronger the objective reason grounded in DI, the stronger is the corresponding subjective reason grounded in Paulo's evidence for DI. And likewise for H.

So, should Paulo swing the bat? Plausibly, Dolores' interest in avoiding harm is greater than Paulo's interest in fulfilling a rather trivial desire to hit the piñata. If so, then the objective reason grounded in DI is stronger than the one grounded in H. So, if DI and H are both true and no other reasons are in play, Paulo will have an all-things-considered objective reason to refrain from swinging the bat. But depending on his evidence for DI and H, he may anyway have an all-things-considered *subjective* reason *to* swing the bat. Perhaps he can make out Dolores' voice coming from a distant corner of the yard, for example, or perhaps his friends are assuring him that everyone is at a safe distance. That would bring the epistemic probability of DI closer to zero, thus weakening Paulo's subjective reason to refrain from swinging, and making it more likely that his reason *to* swing will outweigh it, especially given his strong evidence for H. But even if Paulo's subjective reason to refrain from swinging is outweighed, it will still remain in force, always counting against the act to at least some degree. That makes it a *pro tanto* reason. Thus it follows that Paulo has a *pro tanto*, subjective, victim-affecting moral reason to refrain from swinging the bat.

I want to emphasize two final points about this example before moving on to a case involving killing. First, because Paulo's reason is subjective, it does not depend on the truth value of DI. All that matters is his evidence. As long as he is unable to rule out DI, he has a subjective moral reason not to swing the bat. And since his inability to rule out DI is consistent with DI's falsity, it follows that his reason not to swing the bat is also consistent with DI's falsity. Second, because Paulo's reason is grounded in Dolores' interest in avoiding harm, and because she has this interest regardless of DI's truth value, Paulo's reason is victim-affecting regardless of DI's truth value. This is important because it means that Dolores' interests, together with Paulo's evidence, can give Paulo a reason not to swing the bat even if DI is false—that is, even if swinging it would not harm Dolores. The lesson, as it relates to this paper's argument, is crucial: just because an action will not harm someone does not mean that their interests give us no reason to refrain from doing it.

We can now apply these ideas to a case involving killing. Suppose that Killer is about to kill Victim. Here is a proposition: Epicureanism is false. Call this proposition ~E. As per premise (1), if Killer knew ~E to be true, then he would have a *pro tanto*, subjective, victim-affecting moral reason not to kill Victim. And Killer has at least some reason for thinking that ~E *is* true. Even if he is a dyed-in-the-wool Epicurean, he has no conclusive proof of that view.[12] So the epistemic probability of ~E, given Killer's evidence, is greater than zero. Killer therefore has a *pro tanto*, subjective, victim-affecting moral reason not to kill Victim: *pro tanto* because

[12] And even if he did have such a proof, he might still be unable to guarantee that it really is conclusive.

it might be outweighed, subjective because it depends on Killer's evidence,[13] and victim-affecting and moral because it is grounded in Victim's interest in avoiding harm rather than Killer's interests or desires. And since each of these groundings is logically independent of ~E, so is Killer's reason.

Crucially, Killer's reason is victim-affecting even if ~E turns out to be false. That is, even if Victim is not harmed by being killed, his interest in avoiding harm can still ground Killer's reason to refrain from killing. Remember the above lesson: Just because an action will not harm someone does not mean that their interests give us no reason to refrain from doing it. At worst, the falsity of ~E means that Victim's interests cannot ground an *objective* reason not to kill him, because the actual circumstances will be such that killing him does not damage his interests. But subjective reasons are not affected in this way, because actual circumstances are irrelevant to them. They depend solely on our evidence—in this case, Killer's evidence. And that evidence is compatible with both the truth and falsity of ~E.

I conclude that Killer has a victim-affecting moral reason to refrain from killing Victim, and that this reason applies—and is victim-affecting—even if Epicureanism is true. Just as Dolores' interest in avoiding harm gives Paulo a reason, given his evidence, to refrain from swinging his bat even if doing so will not harm Dolores, so Victim's interest in avoiding harm gives Killer a reason, given Killer's evidence, to refrain from killing, even if doing so will not harm Victim. And since there is nothing special about the case of Killer and Victim, this conclusion is easily generalized: Since none of us is in a position to rule out ~E, each of us has a *pro tanto*, subjective, victim-affecting moral reason not to kill anyone—a reason which applies and is victim-affecting even if Epicureanism is true.

If this is correct, then we are just a short distance away from showing that Epicureanism is consistent with grounding the wrongness of killing in the interests of the victim. All we need to do is move from talk of subjective moral reasons to talk of subjective moral wrongness. And that is easily done: An action (say, killing Victim) is *prima facie* subjectively wrong if and only if we have a *pro tanto* subjective moral reason to refrain from doing it. And when our reason to refrain is an all-things-considered subjective moral reason, we can drop the *prima facie* qualifier. That is, an action is all-things-considered subjectively morally wrong if and only if we have an all-things considered subjective moral reason to refrain from doing it. Thus, in showing that Epicureanism is consistent with a subjective, victim-affecting moral reason not to kill, the above argument establishes that Epicureanism is consistent with a victim-affecting account of the wrongness of killing.

4 Two Objections

My remaining aim is to contrast this account with some recent rivals. But first I want to address two objections.

[13] In particular, it depends on the failure of Killer's evidence to rule out ~E.

The first objection is that my conclusion is too weak to be interesting. The reason that it claims we have is only *pro tanto*, and until we know more about how strong that reason is compared to others that might be in play, we cannot rule out that it will be too easily outweighed. And if it is too easily outweighed, then the ethics of killing that emerges from it will not match our convictions in practice. What we really need, the objector will say, is not just *any* victim-affecting reason to refrain from killing, but one which is strong enough to trump competing reasons in those cases where we judge killing to be all-things-considered wrong. And I have not shown that we have a reason of that strength.

Indeed I have not. But the reason I have argued for is compatible with *any* view about how *pro tanto* reasons to refrain from killing stack up against competing reasons in favor. Sure, that includes views according to which such reasons are easily or always outweighed, but it also includes ones according to which they are difficult or even impossible to outweigh. The ease or difficulty with which the reason I have argued for is outweighed therefore depends on which of these many views is true. Now, the objector is correct that, until we know which one of them is true, we cannot rule out that my reason will be too easily outweighed. But that is true of any *pro tanto* reason to refrain from killing—and, for that matter, any *pro tanto* reason at all. So the present objection does not point to anything that is problematic about my account in particular.[14]

More importantly, the present objection is guilty of moving the goalposts. The initial worry about Epicureanism was that it is inconsistent with victim-affecting reasons to refrain from killing. I have now argued otherwise. In raising a concern about the strength of the reason I have argued for, the present objection effectively concedes my point. That is, it concedes that consistency is not the issue—that Epicureanism is indeed compatible with a victim-affecting reason not to kill. It merely worries that this reason may be too weak to match our moral intuitions. But the very fact that we are now worried about strength rather than consistency is already a sign of important philosophical progress. Indeed, it is a sign of all the progress that I intended this paper to make.

The second objection criticizes my account for making the wrongness of killing depend upon the killer's evidence. Because different people have different evidence, they will assign different epistemic probabilities to ~E. The strength of their *pro tanto* reasons to refrain from killing will therefore vary: those with more evidence for ~E will have a stronger reason to refrain from killing than those with less. This makes it possible for the all-things-considered wrongness of killing a person to vary between two killers merely because of a difference in their evidence. For example, it may turn out that it is wrong for Killer to kill Victim, but not for his brother, Slayer, to do so, because Slayer's evidence assigns a lower epistemic probability to ~E than Killer's does. And this, it might be claimed, is an unacceptable implication.

[14] It is worth adding that, if a certain view about how *pro tanto* reasons stack up against each other implies that the reason I have argued for is too easily outweighed, then we still need some reason for thinking that the fault lies with the reason I have argued for rather than with the relevant view.

But this implication is not unacceptable. After all, the kind of wrongness at issue here is subjective wrongness, and it is entirely proper for subjective wrongness to depend on the evidence of the relevant moral agent. After all, what it is subjectively wrong for us to do is just a function of the subjective reasons we have. And clearly that depends upon our evidence. So, unless it is objectionable for subjective reasons to vary between agents based on differences in their evidence, there can be no objection to similar variations in subjective wrongness.

Furthermore, my argument relies on the claim that, if knowing ~E would give us a reason not to kill, then, to the extent that we have reason to believe ~E, we have a reason not to kill. And although it is plausible that the strength of this reason not to kill should vary according to the strength of our evidence for ~E, it is not strictly necessary. For the reasons I laid out in a previous footnote (n. 8), there is logical space to accept the claim on which my argument relies and at the same time deny that those with more evidence for ~E have a stronger reason to refrain from killing. In short, my argument does not actually produce the implication that the objector alleges. At worst, that implication is produced by my argument *combined with* a certain assumption about how the strength of a subjective reason is affected by the weight of an agent's evidence. And although I do find this assumption plausible, it is logically independent of the case I am making. We can reject it without harm to my argument.

Lastly, as with the first objection, this one effectively concedes the main point of my argument. It worries about the implications of the reason I have argued for, but does not dispute its consistency with Epicureanism. Again, that is already significant philosophical progress. Consistency is all that I intended to prove.

5 Advantages of This Account over Two Rivals

Two other attempts to reconcile Epicureanism with the wrongness of killing have been offered recently, by David Hershenov and Mikel Burley, respectively. In what remains of this paper, I discuss some advantages that my account enjoys over theirs.

5.1 Hershenov's Account

On Hershenov's account, killing is wrong when and because it prevents the victim from enjoying goods that she would have enjoyed if she had not been killed. One instantly wonders how this could be compatible with Epicureanism. If killing can prevent one from enjoying goods, surely death can too. Indeed, if killing prevents the enjoyment of goods, that can only be because it causes death. And if death can prevent the enjoyment of goods, isn't that something bad about it?

Hershenov answers by distinguishing between death, the event, and the state of being dead. He then claims that Epicureanism denies only the badness of the latter (Hershenov 2007: 175):

> Death can mean the event between dying and being dead that the doctor declares at say the time of 10:33 p.m. or it can refer to the period that fol-

> lows that event during which one is dead. A widow might be disturbed by her husband's death because she keeps visualizing the gory end that befell him when a truck hit him as he crossed the street or she may be referring to her lonely days and nights in a big house that feels empty without him. What must be compared to meet Epicurus' challenge is being dead to being alive.

Hershenov claims that his account is compatible with Epicureanism because it makes no such comparison: It does not say that being dead is harmful or worse than being alive, but just that, where the additional years of life would have been worth living, killing is wrong in virtue of preventing the victim from benefitting (Hershenov 2007: 177). These benefits are consistent with Epicureanism because, unlike the badness of being dead, there is no problem locating them at times. We benefit from good lives precisely when we live them. And since the benefits of being alive are compatible with Epicureanism, any account of the wrongness of killing grounded in those benefits will also be compatible with it. That, at any rate, is how I understand Hershenov's position.

I doubt that this approach is successful. Again, if killing can prevent good things, then that is only because it causes death and death prevents good things. The same is true of death and being dead: If death can prevent good things from accruing to the deceased, then that can only be because it causes the deceased to be dead. For example, if my death today can prevent me from enjoying a party on Saturday, then that can only be because it causes me to be dead on Saturday, and I cannot enjoy a party if I am dead. So the ability of death and killing to prevent good things depends on their causal link to being dead: the latter is the only member of the trio that can prevent things non-derivatively, that is, in its own right, and not just in virtue of causing something else that prevents them.

Here is the point: Hershenov's account works only if killing prevents the victim from enjoying good things. But killing can do this only if being dead can do it. And Hershenov's Epicureanism denies that being dead is in any way bad, or worse than being alive. So either being dead cannot prevent the deceased from enjoying good things after all, or, if it can, then that does not count against it (otherwise there *would* be something bad about being dead). Neither possibility bodes well for Hershenov's account. If being dead cannot prevent the enjoyment of goods, then neither can killing. In that case, Hershenov's grounds for condemning killing disappear. On the other hand, if being dead can prevent good things but this does not count against it, then it should not count against killing either. And if it does not count against killing, then how can it serve to make killing wrong?

We can even go a step further. The standard answer to why death is bad for the deceased is that it deprives her of goods she would have enjoyed if she had not died when she did (given that death can deprive only if being dead can, the same answer can be reformulated in terms of the latter). Now, as Hershenov recognizes, the natural answer to *when* deprivations are bad for their victims is that they are bad during the times at which the victims are deprived—the times at which they lack the relevant goods (Hershenov 2007: 173–174). In the case of death, these are times at which the victim is dead. So, if death is going to be bad for people, "then it surely

must be bad for [them] when they are dead" (Hershenov 2007: 174). So why can it not be bad for them at those times? According to Hershenov, it is because being deprived "means one must be in a deprived or harmed state. But the dead are not in any state at all" (*ibid.*). Notice what this implies: Hershenov is not denying that the deprivation associated with being dead is bad; he is denying that there is such a deprivation in the first place. There cannot be one, he says, because that would require the dead to be in a deprived state.

This rationale finds nothing but agreement from me, but it spells trouble for Hershenov's account of the wrongness of killing. As far as I can see, what holds for deprivations holds also for preventions. If so, then Hershenov's reason for denying that the dead can be deprived implies that they cannot be prevented from benefitting either: "To be prevented is to be in a prevented state," he might as well have said. "But the dead are not in any state at all." Now, since no killing occurs unless the victim dies, this means that no killing occurs unless the victim is rendered stateless. And since being prevented from benefitting requires one to be in a prevented state, this means that no killing occurs unless the victim is rendered incapable of being prevented from benefitting. Simply put: If killing causes the victim to become stateless, then it cannot be condemned on the grounds that it puts her into a prevented state. Hershenov's account of the wrongness of killing is inconsistent with his Epicureanism.

The fundamental problem is that he grounds the wrongness of killing in *damage to* the victim's interests, specifically the interest in living a longer life. But since any damage to that interest as a result of being killed can only be due to the fact that killing renders the victim dead, this sort of grounding is bound to conflict with Epicureanism. My account has an advantage here, because it grounds the wrongness of killing not in damage to the victim's interests but in the interests themselves, regardless of whether they are damaged by killing. All that matters on my account are the following two things: First, that if Killer knew that Victim's interests would be damaged by killing, then this would give Killer a reason not to kill; second, that Killer is unable to rule out that Victim's interests will, in fact, be damaged by killing. And since both of these things can hold independently of whether or not Victim's interests are actually damaged by killing, neither of them is in any tension with Epicureanism.

In short, my approach can accomplish what Hershenov's cannot: It can ground the wrongness of killing in the victim's interests without implying that killing damages those interests.

5.2 Burley's Account

The second account I want to look at is that of Mikel Burley. Burley suggests that Epicureans could regard the wrongness of killing as a "basic moral certainty"—a conviction we are entitled to even if we have no rationale to back it up (Burley 2010: 73–79). To illustrate the idea, he quotes approvingly from Jeff McMahan:

> For all of us there are certain convictions which constitute more or less fixed points in the system of our moral beliefs. When a conviction of this sort

> clashes with the dictates of some moral theory, the challenge from the theory must be more rationally compelling than the conviction itself if the conviction is to be justifiably dislodged. Theories seldom satisfy this demand (McMahan 1988: 39).

Burley's idea is that the wrongness of killing is a moral conviction of this sort. Since it requires no rationale to support it, it is not threatened by the truth of Epicureanism. After all, the assumption driving the objection from killing is that the badness of death is a necessary condition for the wrongness of killing. But if its wrongness is a basic moral certainty needing no rationale to justify it, then it has no necessary conditions for Epicureanism to undercut. Killing can therefore be as wrong as ever, even if death is not in any way bad for the one who dies.

I suppose that most people probably do regard the wrongness of killing as a basic moral certainty—a fixed point that moral theories fail to accommodate at their own peril. But then again, most people probably think that about the badness of death as well (indeed, it is in that context that McMahan makes the statement just quoted). But that does not stop Epicureans from denying that death is bad. Nor should it. The truth is not determined by vote, and the fact that most people hold certain convictions does not give those convictions a free pass on justification. Epicureans are well within their rights in claiming not only that the badness of death needs justification, but also that the justifications on offer are no good. Whether or not they are right about this is hotly contested, but the mere fact that Epicureanism is radical by the standards of common sense is no objection to it. Nevertheless, eyebrows begin to rise when the defenders of such a radical view attempt to parry an objection from killing by morphing themselves into fist-pounding moral conservatives. To claim that the wrongness of killing is a certainty needing no justification while the badness of death does need justification may not be contradictory, but it sure yields an odd combination of revisionism (about death) and traditionalism (about killing) that is intellectually very unsatisfying.

None of this yet shows that Burley's approach is unsuccessful, but here is an argument against it. If the wrongness of killing is to have any chance of being a basic moral certainty, then it has to be restricted to a particular class of beings, like sentient beings or persons. For it is clearly not a basic moral certainty that it is wrong (even *prima facie*) to kill dust mites, plants, bacteria, or the skin cells on one's nose, for example. If anything, it is a basic moral certainty that killing these sorts of beings is not wrong (or at least not in virtue of what it does to them). Let us assume, therefore, that the following two things are basic moral certainties: First, that people are *prima facie* wronged by being killed; second, that dust mites are not. I trust that Burley would accept this. But if both of these things are basic moral certainties, then there must be some difference between persons on the one hand and dust mites on the other which explains why persons are wronged by being killed while dust mites are not. Otherwise, our moral system would be inconsistent. It would hold that it is *prima facie* wrong to kill one being but not another, even though there is no morally relevant difference between the two. Now, I take it that, if anything qualifies as a basic moral certainty, it is that no acceptable moral system can be inconsistent.

So, given our two basic moral certainties, there must be some difference between persons on the one hand and dust mites on the other that is morally relevant with respect to killing. And "morally relevant with respect to killing" means that this difference must justify the wrongness of killing persons without also justifying the wrongness of killing dust mites. But if there must be this sort of difference, then it follows that the wrongness of killing persons is not a basic moral certainty after all. Given the permissibility of killing dust mites, it requires justification.

Suppose that this argument fails. Even so, Burley's account requires that the wrongness of killing be a basic moral certainty. My account does not require this, but it is consistent with it, since the fact that something requires no justification is not the same as that it admits of none. Certainly that is an advantage of my account over Burley's.

Moreover, even if it turns out that the wrongness of killing is a basic moral certainty, justifications would still be desirable. The reason is that killing, and even killing persons, may not always be wrong. It might be permissible in cases of voluntary euthanasia, self-defense, or when necessary to save a greater number of others, for example. Whether some killings of persons are permissible is a subject of debate, but at the very least, the wrongness of killing in these sorts of cases is not a basic moral certainty. The very existence of the debate is a testament to that. So if we simply claim that the wrongness of killing persons is a basic moral certainty and leave it at that, then it will be mysterious why killing persons is sometimes permissible, or at least not obviously wrong. Justifications can help to solve that mystery. They can explain the occasional permissibility of killing persons, or at least the occasional lack of certainty about the wrongness of killing them, by pointing out that the considerations in which that wrongness is usually grounded are either absent or count in favor of killing.

Here is an example. Suppose that it is permissible to kill people when doing so is necessary to save a greater number of others. The account I have been defending could explain this as follows: If Epicureanism were false, then the total harm suffered by the larger number of people would exceed that suffered by the smaller number, other things being equal. Since we have (at least a *pro tanto*) reason to minimize harm, this gives us (at least a *pro tanto*) reason to kill the few to save the many. And given our evidence, Epicureanism could be false. Therefore, we in fact have a (*pro tanto*) reason to kill the few to save the many.[15] Now, perhaps this reason does not suffice to make such killing permissible. Perhaps the few have rights against being killed that trump our reason to minimize harm, for example. But even then, the above explanation can at least help to account for why these sorts of situations present such difficult moral questions, and why we find our intuitions pulling us in different directions when we try to answer them. That is already more than what can be accomplished by the mere claim that the wrongness of killing requires no justification.

[15] Hershenov's account could be used to tell a similar story. The reason to kill the few would be grounded in the smaller total of good life that would be prevented by this.

In short, the account defended here is fully compatible with Burley's, but has two advantages over it: First, it does not depend on the wrongness of killing being a basic moral certainty; second, it can go at least some way towards explaining why that wrongness is not always certain, and why killing people may sometimes even be permissible. We can therefore think of my account as not so much a rival to Burley's as an improvement on it—a better execution of the same basic strategy, which is to ground the wrongness of killing in something other than damage to the victim's interests.

6 Conclusion

The primary purpose of this paper was to offer a victim-affecting moral reason to refrain from killing which does not depend on the claim that death is any way bad for the deceased. I argued that we have such a reason because, given our evidence, killing people might harm them, and if we knew that it did, then we would have a victim-affecting moral reason not to kill them. I then argued that, because this reason is grounded in the victim's interest in avoiding harm, and because that interest exists independently of whether killing damages it, it follows that this reason is victim-affecting even if Epicureanism is true. We therefore have a victim-affecting moral reason to refrain from killing which is consistent with Epicureanism about death. The worry that Epicureanism contradicts one of our deepest moral convictions is therefore unfounded.

References

Burley, Mikel. 2010. Epicurus, Death, and the Wrongness of Killing. *Inquiry* 53 (1): 68–86.
Feit, Neil. 2016. Comparative Harm, Creation, and Death. *Utilitas* 28 (2): 136–163.
Hershenov, David B. 2007. A More Palatable Epicureanism. *American Philosophical Quarterly* 44 (2): 171–180.
Huemer, Michael. 2013. An Ontological Proof of Moral Realism. *Social Philosophy and Policy* 30 (1–2): 259–279.
Marquis, Don. 1989. Why Abortion is Immoral. *Journal of Philosophy* 86 (4): 183–202.
McMahan, Jeff. 1988. Death and the Value of Life. *Ethics* 99 (1): 32–61.
McMahan, Jeff. 2002. *The Ethics of Killing: Problems at the Margins of Life*. New York: Oxford University Press.
Pleasants, Nigel. 2008. Wittgenstein, Ethics, and Basic Moral Certainty. *Inquiry* 51 (3): 241–267.
Silverstein, Harry S. 1980. The Evil of Death. *Journal of Philosophy* 77 (7): 401–424.
Yourgrau, Palle. 2000. Can the Dead Really be Buried? *Midwest Studies in Philosophy* 24 (1): 46–68.