

Against Self-Defense

Blake Hereth

Abstract: The ethics of self-defense is dominated by the Orthodox View, which claims that at least some cases of self-defensive assault are permissible. I defend the radical view that there are no permissible instances of self-defensive assault. My argument proceeds as follows: Every permissible act of self-defensive assault could, in principle, have its permissibility be massively overdetermined. Such ‘super-permissible’ acts of assault are ones in which agents are objectively permitted to perform those acts in morally trivializing or cavalier fashion: that is, agents need not ‘think twice’ about inflicting or permitting harm and are permitted to assault persons as if it were morally insignificant. Yet this is never true, since assaulting persons is always morally serious. It follows that there are no acts of permissible self-defensive assault.

Keywords: self-defense, Orthodox View, Cavalier Constraint, Near-Cavalier Constraint, super-permissibility

§1 Introduction

This essay concerns what I shall call the Orthodox View of self-defensive assault:

The Orthodox View: Some instances of self-defensive assault are morally permissible.

By ‘assault,’ I mean the sort of physical or psychological harm characteristically inflicted in self-defense.¹ Here is one useful account:

You assault me at time t only if and because:

- (i) You intentionally and directly cause an event, C ;
- (ii) prior to C , a fully informed me would have had conclusive prudential reason to escape C (absent some other C -type circumstance if I succeeded in escaping C) if I aimed to protect myself; and
- (iii) I do not escape C .

1. For another account, one that characterizes these acts not as acts of assault but as acts of *violence*, see Audi 2009. Audi notes that violence can be done accidentally, but accidents are not the sorts of things there can be a moral obligation to avoid, and the sort of pacifism I defend is principally contrasted with the Orthodox View in virtue of disagreements about the sort of harmful act which is always and everywhere wrong.

On this account, several essential features of assault are present.² The first condition identifies two: *intentionality* and *directness*. Any unintentional harm is not assault; there can be no accidental assault. Assault is therefore purposeful, intentional. Regarding directness: if you hire an assassin to kill me, then *you* do not assault me. At most, you intend my assault. Thus, assault requires causal directness. There are degrees and kinds of directness, however, but the directness condition is designed primarily to exclude as instances of assault cases in which some agent (qua agent) distinct from you harms me.

I am *fully informed*. Consider a case in which, prior to C, I have conclusive prudential reason to escape C, but this conclusive prudential reason is grounded in some misleading epistemic consideration: You aim a gun at me but don't intend to fire. This would hardly constitute assault. Thus, I must (again, prior to C) have conclusive prudential reason to escape C, where this conclusive prudential consideration holds on the assumption that I am fully informed.

The circumstance, C, must itself be *counterfactually sufficient to motivate conclusive prudential reason to avoid C*. Consider a case in which I attempt to escape C (say, by running out the back door of the pub) but, when attempting to escape, finds myself in another C-type circumstance (say, there is a violent gang outside ready to mutilate me). Staying in the bar and being shot in the leg might, from my standpoint, be prudentially *better* than attempting to escape being shot in the leg.³ Still, this would not make shooting me in the leg a case of non-assault. Circumstance C is one for which I have conclusive prudential reason to escape absent some alternate C-type circumstance.

Escaping C must be for the sake of *me*. Someone might have conclusive prudential reason to move their car from the train tracks because, otherwise, someone else will intentionally and directly smash it to bits. But this would not be assaulting me because it is not an act against me, even if it might exert causal effects on me.

Finally, that which I have conclusive prudential reason to avoid must *happen*. After all, I might have conclusive prudential reason to avoid C (where C is shooting me in the leg); there might be no other pressing C-type circumstance (no gang outside); and still I might escape C. Upon doing so, I will have avoided assault while meeting the other conditions. Thus, the prior conditions

-
2. I exclude force as a requirement, since it is not obvious to me that force is a necessary condition of assault.
 3. Perhaps this is why consensual euthanasia does not seem to be a case of assault. If it is in the person's best interest to die and if the person consents, it seems to many that killing the person (or, on passive variations, intentionally letting the person die) is compassionate and hardly an instance of providing prudential reason to avoid C. Where it is best for someone not to escape but rather to die, it is not an act of assault to kill her (or, on passive variations, intentionally to let her die).

are not sufficient for assault. It is necessary that I fail to escape C for the account to be completed.

Widespread acceptance is enjoyed by the thesis that self-defensive assault of this kind is sometimes morally permissible. Thomson (1991: 283) says that many “think of self-defense as morally transparent. What could be clearer than that morality permits a person to save his or her life against threats to it?” She then asks her readers to consider the case of Villainous Aggressor:

[Y]ou are standing in a meadow, innocently minding your own business, and a truck suddenly heads toward you. You try to sidestep the truck, but it turns as you turn. Now you can see the driver: he is a man you know has long hated you. What to do? You cannot outrun the truck. Fortunately, this is not a pure nightmare: you just happen to have an antitank gun with you, and can blow up the truck. Of course, if you do this you will kill the driver, but that does not matter: it is morally permissible for you to blow up the truck, driver and all, in defense of your life. (Thomson 1991: 283)

Thomson is explicit that this moral permission is not mere excuse, but that blowing up the truck (and the driver) is morally permissible, justified killing.⁴ This is because

blowing up the truck in Villainous Aggressor is not something you ought not do. We cannot plausibly say that you ought not blow up the truck, but will only be in a measure at fault, or in no measure at fault, for doing so: you simply *may* blow up the truck. Morality permits it. (Thomson 1991: 283)

Cases like Villainous Aggressor are widely regarded as paradigmatic for justified instances of self-defensive assault. Jonathan Quong explicitly endorses Thomson’s conclusion about Villainous Aggressor and other cases Thomson raises, writing, “Like Judith Jarvis Thomson, I think these are all cases where it is permissible to kill one person in order to save my own life” (Quong 2009: 507). Quong is hardly alone in this endorsement. Many philosophers endorse Thomson’s conclusion about Villainous Aggressor and structurally similar cases, and in their comments suggest consensus:

In these circumstances, few of us would condemn you for killing in self-defense. Nor would we condemn a third party who intervened on your behalf by killing your attacker. (Draper 1993: 73)

Draper’s talk of condemnation appears to support the view that defenders are in many cases unworthy of condemnation because they act rightly.

[T]he simple point is that if somebody else is intent on leaving you dead, and there is strong immediate evidence that the only really feasible way to

4. Thomson’s claim that “it does not matter” if the driver is killed is perhaps best understood as: it does not matter, all things considered (which is compatible with, but does not entail, the stronger claim that the death of the driver is altogether morally irrelevant).

prevent this is to leave him dead first, then normal people prefer the latter to the former. (Narveson 2003: 158)⁵

It is widely accepted, for example, that if one person, Albert, culpably threatens the life of another person, Betty, then he may have forfeited his rights against serious harm being imposed on him in Betty's defence. (Quong and Firth 2012: 674)

Still others, such as Jeff McMahan (1994a: 193),⁶ Michael Otsuka (1994: 74), and Helen Frowe (2014: 1) claim that this view is widely held. If these judgments about Villainous Aggressor cases are true, then self-defensive assault is at least sometimes justified.⁷ Beyond agreement about Villainous Aggressor cases, there is even broader agreement that self-defense is permissible. Draper (2009: 69) refers to self-defense as "perhaps the most widely recognized justification for inflicting harm." Tyler Doggett (2011: 220) makes the broad claim that "[t]he literature on self-defense agrees that killing" Villainous Aggressor is permissible. Quong (2012: 45) endorses the Orthodox View unequivocally: "Sometimes it is morally permissible to seriously harm or kill people in self-defense, or in defense of others."

Some philosophers even contend that the permissibility of self-defensive assault in some cases is either itself obvious or is itself entailed by obviously true moral principles. Among them are Seth Lazar (2009: 700), McMahan, F. M. Kamm (2012: 219), and David Rodin:

Defensive rights seem to be entailed in a very basic way by rights to things. Thus if I have a right to X, then it seems to follow as a simple corollary that I have the right to take measures to prevent my right to X from being violated. (Rodin 2002: 37)

McMahan (1994c: 252) goes so far as to say that the Orthodox View is an axiom of contemporary ethical theory. Concurring with McMahan that ordinary morality includes the possibility of justified self-defensive assault, Lazar (2009: 700) claims that self-defense can remove "*the whole wrong* in killing a person" (emphasis mine). This suggestion makes theories of self-defensive attractive tools for those who believe that war can be justifiably waged. Rodin's powerfully simple observation that a right to something entails a right to preserve it supports the permissibility of self-defensive assault: if one's right to life is (wrongly) threatened, one is permitted to exercise one's right to do what is necessary to protect that right, including in cases where killing or seriously harming others is necessary to protect one's right to life.

5. Narveson's use of "normal people" is suggestive of consensus.

6. McMahan goes further than some in suggesting that self-defensive assault is just as often obligatory as it is merely permissible. Still, McMahan is a clear instance of a philosopher who endorses the permissibility of self-defensive assault.

7. Draper and Frowe make similar judgments: see Draper 1993: 73 and Frowe 2010: 247.

I shall argue that the Orthodox View is false. My argument will proceed as follows. If it's ever permissible to harm an individual with the sort of moral status most human beings are believed to have, that permissibility can be overdetermined. Furthermore, it can be massively overdetermined. Thus, every permissible act of self-defensive assault has the modal property *being possibly massively overdetermined*. Where the permissibility of harming a being is massively overdetermined, I say it is "super-permissible." I then argue super-permissible acts of assault are ones in which agents are objectively permitted to perform those acts in morally trivializing or cavalier fashion. Among other things, this means that agents need not 'think twice' about inflicting or permitting harm; they are permitted to assault persons as if it were, morally speaking, insignificant.⁸ Yet this is never true, since assaulting others is always deeply morally serious. It follows that there are no acts of permissible self-defensive assault. Though the focus of this paper will address *self-defensive* assault, it should be reasonably clear that a similar argument could be given against other kinds of assault (e.g., other-defensive or punitive).

§2 Permissible Action and the Cavalier Constraint

I claim that there is a *Cavalier Constraint* for permissible action. This constraint is a side-constraint in the sense described by Robert Nozick: an act is permissible for some agent only if it does not violate any (absolute) side-constraint.⁹ This constraint is described as follows:

-
8. My argument is therefore distinct from Andrew Fiala's (2014b) argument against the Orthodox View. Fiala's argument is that the Orthodox View gets its moral support by avoiding hard cases and focusing instead on sterilized 'Trolley' cases in which it appears obvious that self-defensive assault is justified. Because such cases abstract away from more difficult cases, Fiala claims the Orthodox View gets its moral support too cheaply. Fiala's argument is an important one, but it's not the one I want to make. The sense in which killing—or, to use my favored term, 'assault'—becomes problematically *cavalier* is the sense in which its performer would be objectively justified in acting as if assaulting someone was no big deal, given the host of considerations that make the act permissible. My worry is that the Orthodox View implies that this is possible, which I use as a *reductio ad absurdum* against the Orthodox View. There are, then, two possible approaches here: The first is Fiala's, which points out that killing people often gets messier than we expect, and we can never assume that killing is as sterilized as 'Trolley' cases make it appear. The second is that regardless of how things *actually* go when we kill people, the Orthodox View implies that we can kill people as if it's no big deal, and that's false. The difference between Fiala's argument and mine is that Fiala thinks defenders of the Orthodox View get more than they bargain for even if they're right about the sterilized cases, whereas I think they're wrong about even the sterilized cases. My thanks to an anonymous reviewer for noting the similarities between my argument and Fiala's and for pressing my to clarify how my argument is distinct from his.
9. See Nozick 1974: 28–33.

Cavalier Constraint: There is some class of actions that can't be permissibly treated cavalierly, where an action ϕ is treated cavalierly if there is some possible world in which some agent is objectively morally permitted to treat or regard the doing of ϕ as if it were morally "no big deal."

What I shall show in this section is that the following premise is true:

Premise 1: The Orthodox View is true only if there is some instance of permissible self-defense that could not, in principle, violate the Cavalier Constraint.¹⁰

Imagine that a moral agent comes permissibly to regard the performance of some action as if it were morally inconsequential. The sense in which this agent acts permissibly is objective, and not (merely) epistemic:¹¹ Her action is made permissible by some moral consideration, and she knows this. Suppose that she regards ϕ -ing in a morally cavalier way by treating it with a sort of moral ease: She very casually does ϕ and doesn't give a second thought to doing so.¹²

The Cavalier Constraint doesn't apply to *all* actions. It remains open that there are actions that *can* be done cavalierly. For example, suppose ϕ -ing is the act of respecting the wishes of my friend to leave her favorite ceramic penguin facing North, but imagine that someone has filled the penguin with explosives that will kill my friend unless I knock it off the shelf, destroying it and averting the blast. In such a case, failing to ϕ is not only *permissible*; it's so very permissible that I would be objectively justified in knocking the ceramic penguin off the shelf without a second thought.

Whereas some actions can be permissibly done in cavalier fashion, however, others can't. To use just one example of a morally grave action, consider torture. Let ϕ be the act of not torturing someone and assume, for the sake of

10. Said another way: If the Orthodox View is true, then there is no instance of permissible self-defense that could, in principle, violate the Cavalier Constraint.

11. Cf. Jeff McMahan's overview of the distinction (2011: chap. 4, esp. pp. 162–63).

12. "Regarding" might appear tacitly to endorse some form of doxastic voluntarism, since it requires agents neither to treat nor regard a violation or disregarding of ϕ as morally inconsequential. But the sense of regarding need not be doxastic, and the only relevant doxastic cases are ones in which one does have even some voluntary control over one's beliefs. For example, if one voluntarily selects actions one knows will result in believing, or believing more strongly, that ϕ is morally inconsequential, then one would act wrongly in choosing those actions.

Cf. Nagel 2005. Nagel suggests that citizens have an interest in having political leaders who are ruthless and carry out the business of the state, but their ruthlessness should not become internalized. (In the terms employed here, ruthless actions should not be carried out with ease.)

argument, that acting *against* ϕ can be permissible, perhaps on grounds of self-defense or other-defense, as Uwe Steinhoff argues:

There is no morally relevant difference between self-defensive killing of a culpable aggressor and torturing someone who is culpable for a deadly threat that can be averted only by torturing him. (Steinhoff 2006: 337)¹³

In some theories of self-defense, it's permissible to assault unjust aggressors when necessary for self-protection because those aggressors have, via their unjust aggression, *forfeited* their right not to be assaulted. Though many advocates of self-defense take issue with Steinhoff's claim that self-defensive assault is morally indistinct from self-defensive torture, assume that self-defensive torture is sometimes permissible.¹⁴ Consider the following case:

Uneasy Kai knows that self-defensive torture is permissible and finds herself in a circumstance in which torturing a person is their only means of self-defense. They torture the threat but without moral ease: They do not enjoy it, they're cautious to inflict only the harms they deem necessary, and participate with a high degree of hesitation and trepidation.

Uneasy Kai does not enjoy torture in any way; they regard it as unfortunate that it is at all instrumentally necessary for their own defense (or the defense of others). *Kai* realizes too that torture, even in self-defense, is a very morally serious activity, and they responds appropriately by taking great pains to be sure that each cut and each blow is absolutely necessary, giving their threat ample time between each clipped nail to reconsider sharing information relevant to the threat against *Uneasy Kai*.¹⁵ But consider a variation on *Uneasy Kai*:

Nonchalant Kai learns that *Bibi* is nearby, and they decide to kidnap *Bibi* and torture him for information necessary to save *Kai*'s life. With each electrical shock and pried fingernail, *nonchalant Kai* proceeds with moral nonchalance: They don't give much additional thought to whether some harm they inflict is necessary or whether their torturous acts are right.

Unlike *Uneasy Kai*, who responded appropriately to the moral graveness of torture, *Nonchalant Kai* fails to respond appropriately. Torture is intrinsically grave and the moral risks associated with it are plentiful: inappropriate emotional attitudes towards or reactions to harming others; objectifying victims;

13. Cf. Steinhoff 2012: 19–26 and Steinhoff 2013: chap. 9.

14. See the challenge to this conclusion by Bufacchi and Arrigo (2006).

15. This does not assume that there are no metaphysically vague cases in which it is indeterminate whether inflicting some harm *H* is necessary (or, is among a set of comparable alternatives, one of which is necessary) to achieve goal *G*, or that there are no cases in which it is epistemically unclear whether *H* is necessary to achieve *G*. Rather, the claim is simply that no agent is possibly objectively justified in casually assuming that *H* is necessary to achieve *G* when *H* is morally serious.

resorting to inappropriate means (torturing non-threatening family, for example); torturing victims even after they have surrendered information due to uncertainty about the truthfulness of their claim to have surrendered the correct information; and more. Thus, even if torturing in self-defense is permissible, this does not permit anyone to treat torture in any way that would constitute a nonchalant or casual attitude toward its moral seriousness. Everyone everywhere always has a duty to take torture seriously, even if they are inflicting it with justification.¹⁶

Examples like these help to raise two important questions. The first is what makes actions, or action-kinds, subject to the Cavalier Constraint? Second, if those actions or action-kinds cannot permissibly be treated cavalierly, what moral relations and functions (if any) are affected in virtue of the Cavalier Constraint?

It's not necessary for the purposes of this paper to sketch the full set of necessary and sufficient conditions for what qualifies an action for the Cavalier Constraint. It's enough to show that there is a Cavalier Constraint and that it applies to assaulting persons. To this end, one sufficient condition will be identified and defended.

The short defense is as follows. All persons have properties, some of which are essential. One essential property of persons is *being extremely morally considerable*. This property necessarily confers a right to all persons, a right they cannot forfeit and a right that cannot be countervailed: the right not to be treated carelessly or, when harming them, harming them casually.¹⁷

16. Cheyney Ryan's pacifism is of the skeptical variety. He maintains that "the proponent of killing cannot produce a single argument for why killing another person is permissible" (1983: 509). He proceeds to argue as follows: There is no plausible sense to be made either of the claim that Aggressor forfeits his right to life or of the claim that Victim has a right to kill Aggressor, yet one of those is necessary if a defense of the Orthodox View is to succeed. Thus, the Orthodox View doesn't succeed. Where my argument bears some similarity to Ryan's is where he discusses the positive reasons for being a pacifist, namely, "The pacifist's problem is that he cannot create, or does not wish to create, the necessary distance between himself and another to make the act of killing possible" (1983: 521). Ryan offers the example of shooting a fascist with his trousers undone: It doesn't seem at all wrong to kill fascists, but it does seem wrong to kill blokes with their trousers down. What the pacifist refuses to do, therefore, is to come to see unjust aggressors as—to borrow a phrase from James Kellenberger—"conscious, knowing obstacles . . . whose removal or neutralization is a means to one's end" (1987: 136). However, unlike Ryan, my point is not that we should refuse to distance ourselves from others to the extent that we can, psychologically, see them as things to be killed. Rather, my point is that we should reject as false any view that would objectively permit us to distance ourselves from others to the extent that we can kill them as if it were no big deal. I argue that the Orthodox View does so, and that the Orthodox View is therefore false.

17. This is true only if there are rights. It is assumed that there are, but the following judgments about what is permissible and what is not do not depend upon the existence of rights.

The longer defense begins with a defense of the claim that all persons possess that essential property. To begin, consider that every person has a presumptive right against being harmed or killed. This right is absolutely presumptive: it holds absent any forfeiting or countervailing conditions.¹⁸ Said another way, the right against being harmed or killed holds whenever the person being harmed or killed neither forfeits that right nor has that right countervailed. Consider such a case. Bailey is not attacking anyone and there is no moral good worth preserving or evil worth preventing. Thus, there is no moral reason to kill Bailey. If you harm or kill Bailey, you have wronged her, and this is true in all possible worlds in which you harm or kill Bailey under these conditions. By modal implication, Bailey *necessarily* has the right not to be harmed or killed absent forfeiture or countervailing conditions.

As it concerns harming or killing persons, the threshold at which countervailing conditions obtain is high. If a gallon of milk will be spilt unless Bailey is harmed or killed, we ought not to harm or kill her. If Bailey will die unless she tosses Shana under the Greyhound bus, Bailey still ought not to toss Shana under the Greyhound bus. Restrictions on permissible harming and killing are plentiful, even in cases of permissible self-defense.

Moreover, even in cases where countervailing (but not forfeiting) conditions do obtain, it is sometimes the case that there exists some *residual duty* to the person or persons harmed or killed. McMahan, for example, writes:

When one thus permissibly acts against a right, I will say that one *infringes* that right, whereas when one impermissibly does what another has a right that one not do, one *violates* that right. Even though an agent acts permissibly in infringing a right, the victim is nonetheless wronged and may thus be owed compensation. (McMahan 2011: 10)

A paradigmatic example of residual duty is a modified trolley case in which either Steve will lose his leg or five people will die. Assuming it is permissible to switch the track such that Steve loses his leg, it appears true both that Steve had a right not to have his leg taken *and* Steve is owed compensation. When non-combatants are unintentionally harmed as a direct result of tactical bombing or other means of offensive warfare, they too are plausibly owed compensation if such compensation can be provided.¹⁹ More strongly, every person whose right is permissibly infringed in McMahan's sense is owed compensation, other things being equal.

The moral significance of entities like persons, each of whom possesses a necessary presumptive right not to be harmed and a presumptive right to com-

18. Forfeiture conditions operate in the following way: if some circumstance obtains, then person Steve forfeits Steve's right. Countervailing conditions operate in the following way: if some circumstance obtains and is stronger than the relevant right, then Steve's right is permissibly overridden (or trumped).

19. See Brummer 1996.

pensation in cases of right-infringement is nearly inestimable. These rights appear to confer an extremely considerable moral status: the rights are possessed presumptively, have a remarkably high countervailing threshold, and are possessed by every person in all possible worlds in which those persons exist. Furthermore, the extremely morally considerable nature of persons is grounded in the *modal* properties possessed, which persons continue to possess even if, for example, their right not to be harmed or killed is forfeited or countervailed in the actual world. Supposing that Bibi's right to life or not to be harmed is forfeited or countervailed in the actual world, it hardly follows that those rights of Bibi's or their presumptive force are forfeited or countervailed *across worlds*. Even if Bibi forfeits his right not to be assaulted today because he aggresses unjustly against Kai, Bibi does not thereby *absolutely* forfeit his right across possible worlds. Rather, Bibi retains his presumptive right not to be assaulted such that when there is no reason to harm Bibi, then harming Bibi is wrong. These essential properties of persons, therefore, ground the following right of persons: the right not to be treated carelessly or, when harming them, harming them casually.

§3 Relations Affected by the Cavalier Constraint

Because of the Cavalier Constraint, certain moral relations are affected. Relations affected by persons with the relevant rights will include, among other things, the *overdetermined* relation. To say that the permissibility of self-defense is overdetermined is to say that there is more than one independent moral consideration whose normative implications entail permissible self-defense. For example, if one's permission to protect one's children implies that one is permitted to self-defend,²⁰ and if one likewise has an independent permission to protect oneself, and if defending oneself is likewise necessary to prevent catastrophe,²¹ then that set of permissions supplies independent moral grounds for permissible self-defense. The permissibility of self-defense is overdetermined because, in each case, ϕ -ing is permissible and it requires self-defense. That is, protecting one's children requires first that I defend myself, as does preventing catastrophe. Thus, the permissibility of ϕ -ing in each case *grounds* the permissibility of self-defense.

There is conclusive reason to believe that all theories of justified self-defense *in principle* permit overdetermination.²² That is, the following premise is true:

20. See Davion 1990: 92, Parfit 2011: 205, and Narveson 1965: 262.

21. See Steinhoff 2006.

22. Excluding, of course, eliminativist theories of self-defense like pacifism, which imply that there is no permission, and therefore no right, to self-defensive assault.

Premise 2: Any possible instance of self-defense could, in principle, have its permissibility massively overdetermined.²³

According to all Orthodox theories of self-defense, self-defense is not absolutely wrong (that is, it is actually permissible in some possible circumstance) and can be motivated (that is, there would be moral enough reason to permit self-defense in the circumstance). Each of these conditions is strictly necessary for overdetermination.²⁴ If self-defense were *absolutely prohibited*, it would necessarily fail to make any act of self-defense permissible or contribute to its permissibility. Thus, the permission to protect one's children or prevent catastrophe could make self-defense permissible or contribute to its permissibility only if self-defense is not absolutely prohibited.

Moreover, it must be the case that self-defense is possibly motivated. This assumption is not built into the conception of self-defense being actually permissible in some circumstance, since self-defense might be permissible even if there is no moral reason to act in self-defense. The criminal who breaks into someone's home in order to steal a leftover blood sample for a harmless science project might use defensive measures to prevent the vandalism, but it is not clear that there is any *moral* reason to prevent this. The person who finds this activity strange but harmless is hardly lacking in moral reasons-responsiveness when she fails to prevent the robbery.

While these observations alone do not conclusively support the conclusion that all theories of self-defense permit overdetermination, they do some work in clearing the way of conceptual obstacles. In particular, they show that two possibilities that necessarily would make overdetermination impossible do not hold for theories of self-defense. This provides some support for the conclusion that overdetermination is (necessarily) possible for theories of self-defense. Fortunately, stronger support can be given.

An extended defense of the possibility of overdetermination includes cases in which overdetermination appears obvious. There are, necessarily, possible cases in which there is more than one moral consideration that normatively implies a permission to self-defend. Returning to our earlier examples, we have the following set of permissions:²⁵

- (a) All guardians have a *pro tanto* permission to protect children under their care;

23. A reviewer asks what I mean when I say that some act is "massively" overdetermined. By this, I simply mean that the moral considerations that (supposedly) make the act permissible are large in number or have a strong moral valence (or both).

24. Both are likewise necessary for moral determination (again, in the grounding sense).

25. The following set of permissions can easily be substituted for a distinct set where the permissions for self-defense are more obviously such that there is no lexical priority.

- (b) All innocent persons have a *pro tanto* permission to protect themselves; *and*
- (c) All persons have a *pro tanto* permission to prevent catastrophe at serious cost to themselves (but not to others).

Let's also suppose that the following is true at a given time for some agent:

- (d) Unless Kai self-defends right now, Kai will fail to protect children under their care;
- (e) Unless Kai self-defends right now, Kai will fail to protect themselves; *and*
- (f) Unless Kai self-defends right now, Kai will fail to prevent catastrophe.

Suppose also that there are no permissions distinct from these ones which would further imply that Kai was permitted to self-defend right now. Thus, the moral considerations which permit Kai to self-defend are *exhausted* by considerations (a)–(f). Furthermore, assume that there is no lexical or temporal priority amongst the permissions. That is, assume that each permission *has the same moral strength* (in terms of being a moral consideration) and *arises at the same time*. When Kai finds themselves in the circumstance in question, there is no permission that enjoys a 'first-to-ground' privilege in virtue of some lexical or temporal priority. What this means is that no particular permission *uniquely* grounds Kai's permission to self-defend right now. Yet, because Kai's permission to self-defend right now *is* grounded by (a)–(f), it follows that the permission is *overdetermined*.

There is another argument for this conclusion: a best-explanation argument. The permission to self-defend appears stronger in cases in which more is at stake than in cases where less is at stake. If Kai's failure to self-defend would result in the death of their children and the destruction of the greater Baltimore area, then Kai's permission to self-defend is plausibly stronger than in a case in which Kai's failure to self-defend would result merely in their own death. In such a case, Kai's permission to self-defend appears stronger because of the normative influence of certain counterfactual considerations: namely, that the destruction of Baltimore would result and that would be worse. Thus, the fact that the permission to self-defend appears sensitive to these concerns is further evidence for the possibility of overdetermination.

The explanation acquires further strength when it is recognized that the relevant moral considerations (preventing the death of one's family and preventing the destruction of Baltimore) are each individually sufficient to motivate at least a *pro tanto* permission to self-defend. Because of this, each consideration is ordinarily permission-*creating*, and it is difficult to see why it would fail to be permission-*affecting* in these cases.

I shall now argue that overdetermination is ruled out by the Cavalier Constraint. That is, I shall defend the following premise:

Premise 3: Any act that could, in principle, have its permissibility massively overdetermined is also an act whose performer would, under those conditions, be permitted to perform that act in a cavalier fashion.

To see why, recall that assaulting persons *cannot* be permissibly done cavalierly.²⁶ If Kai is to permissibly assault a person, they cannot be acting on the basis of moral considerations that make assaulting the person super-permissible, or so permissible that Kai is permitted to assault the person cavalierly. But if assault can be made permissible and in a way that is overdetermined, a problem arises.²⁷ Suppose that if Kai is to avoid the horrific leveling of Baltimore, they must torture Bibi. As in a previous case, Kai's partner, children, parents, and extended family live in Baltimore. Kai is in Baltimore interrogating Bibi. Grant lastly that torturing Bibi is permissible and overwhelmingly so. If this is true, then Kai's permission to assault Bibi (which includes *torturing* Bibi) is made super-permissible by the relevant set of moral considerations, such that the act assaulting Bibi could be permissibly done cavalierly.

We might tack on a fistful of other permissions that coincide to overdetermine the permissibility of torturing Bibi. In addition to (a)–(f), we might add the following:

- (g) All persons have a *pro tanto* permission to keep their promises;
- (h) All persons have a *pro tanto* permission to demonstrate meaningful opposition to terrorist activities.

We shall also assume, as before, that the following claims are also true:

- (i) Unless Kai self-defends right now, Kai will fail to keep their promises;
and
- (j) Unless Kai self-defends right now, Kai will fail to demonstrate meaningful opposition to terrorist activities.

There might in principle be an infinite number of such permissions, or at least such a swath of them that the permission to assault Bibi is super-permissible. Recall the case in which a friend's wish to have her ceramic penguin face North can permissibly be infringed since failing to move the penguin will otherwise result in loss of life or limb. The savior's *pro tanto* duty to act within the bounds of her friend's wishes is vastly overridden by the permissions

26. It is a necessary truth because, as defended above, it is true in every possible world that if persons exist, they have a presumptive right not to be killed or harmed, which entails a corresponding duty on others not to harm or kill persons absent forfeiture and countervailing considerations. But if some claim is true in every possible world, then it is a necessary truth.

27. The metaphor of the trump card is illustrated powerfully by Dworkin (1984: 153).

grounded in saving lives. It is because of this significantly overdetermined permissibility that the savior would act permissibly in refusing to give even a second thought to tossing the penguin. Something similar could occur in some cases of assault, however, in which self-defensive assault can be and is super-permissible. Yet this conclusion is vastly at odds with the existence of a Cavalier Constraint for assaulting persons. Because there is such a constraint and because the possibility of permissible self-defensive assault entails the possibility of overdetermination, and because such overdetermination would result in trivializing an act for which there is a Cavalier Constraint, we should conclude that *self-defensive assault is not even possibly permissible*.

The argument likewise holds even if one adopts a forfeiture view of permissible self-defense. On a forfeiture view, Bibi *loses* a right, the former presence of which would have entailed a legitimate claim against Kai that Kai not attack Bibi.²⁸ Now Bibi's right is gone, and with it comes the permissibility of killing him. But the permissibility of killing Bibi can still be overdetermined. For example, Frowe (2014: 3) claims that Bibi forfeits his rights through "involuntary behavior," but there might be several simultaneous instances of voluntary behavior by which Bibi forfeits his rights. Perhaps Bibi unjustly attacks Kai and also unjustly attacks a bystander, each of which would be sufficient for Bibi to forfeit his right, rendering permissible an attack by Kai. In that sort of case, Kai is *further permitted* to attack Bibi; it is more important to attack Bibi than it would be if, for example, Bibi attacked only Bibi. Such reasons can be added until the permissibility of assaulting Bibi becomes super permissible, such that Kai would be objectively justified in assaulting Bibi without a second thought. Because this is impossible, given the Cavalier Constraint, and because forfeiture accounts entail that it is possible, forfeiture accounts do not escape the objection.

To complete the argument against the Orthodox View, only one premise is required, and it's true by definition:

Premise 4: Performing an act in a cavalier fashion would violate the Cavalier Constraint.

It follows from this and the three prior premises that the Orthodox View is false.

28. See McMahan 2011: 10. For a defense of the forfeiture account, see Lang 2014 and Frowe 2014: 3–4, 72–73. Notably, Lang (2014: 38) holds that whereas the standard picture of permissible self-defense is one in which the wrongful aggressor loses her right not to be killed, forfeiture accounts are accounts of how that right is lost.

§4 Relations Affected by the Near-Cavalier Constraint

The lynchpin of the argument against the Orthodox View concerns overdetermination. In that argument, it was claimed that any number of distinct moral reasons might normatively affect the permission to self-defend against an aggressor. Because of this, the act of assaulting a person would be trivialized in whichever possible worlds the horde of moral reasons overdetermines the permissibility of assaulting a person. In virtue of this false modal implication, the Orthodox View is false. More formally:

The Argument from the Cavalier Constraint

(Premise 1) The Orthodox View is true only if there is some instance of permissible self-defense that could not, in principle, violate the Cavalier Constraint.

(Premise 2) Any possible instance of self-defense could, in principle, have its permissibility massively overdetermined.

(Premise 3) Any act that could, in principle, have its permissibility massively overdetermined is also an act whose performer would, under those conditions, be permitted to perform that act in a cavalier fashion.

(Premise 4) Performing an act of self-defense in a cavalier fashion would violate the Cavalier Constraint.

(Conclusion) The Orthodox View is false.²⁹

What might be said against this suggestion is that the argument itself provides principled reason to reject the very possibility of *cavalier* overdetermination while maintaining the possibility of *non-cavalier* overdetermination. This might appear ad hoc, but it need not be. The argument offered here might as well be an argument against the possibility of excessive overdetermination, and one might suppose that this is logically compatible with the truth of the Orthodox View.³⁰ Moreover, this suggestion is compatible with agnosticism about what the precise cutoff conditions for normative influence are: which moral considerations are disqualified, to what extent normative influence can occur, and the like. Thus, the objector might concede that the argument above would be compelling if there were not countervailing considerations, such as the objection argument provides. More specifically, the Argument from the

29. A reviewer inquires why a more modest conclusion does not follow (e.g., “The Orthodox View is false, in principle”). As footnote 10 makes clear, the first premise is identical to the conditional claim: “If the Orthodox View is true, then there is no instance of permissible self-defense that could, in principle, violate the Cavalier Constraint.” Moreover, premises (2)–(4) jointly entail the falsity of the consequent of the first premise. Thus, by *modus tollens*, the stronger conclusion follows.

30. Unless, as has been argued here, and as will be further argued in a moment, the Orthodox View does entail the possibility of cavalier overdetermination.

Cavalier Constraint would succeed in showing that cavalier overdetermination is made possible by the Orthodox View *only if* there is no principled way for the Orthodox View to rule out cavalier overdetermination. But that is precisely what is made possible by the argument against the Orthodox View.

Beyond trivializing the act of assaulting persons, there is a secondary constraint for permissible action. Consider the following case:

Almost Nonchalant Kai captures and tortures Bibi to prevent a catastrophe. Kai realizes that torturing Bibi cannot be permissibly done cavalierly, and they do not demonstrate a total lack of moral attentiveness to whether their actions are necessary, right, etc. They are not quite cavalier. But they are *close* to being cavalier: They remain very comfortable, assuming, and regard torturing Bibi as *nearly* inconsequential.

While not acting quite as wrongly as Nonchalant Kai, *Almost Nonchalant Kai* *does* act wrongly. They act wrongly because torturing a person is not near-trivial, and thus it is wrong to treat torturing a person as if it is near-trivial. Thus, for some acts, there is a *secondary* constraint:

Near-Cavalier Constraint: There is some class of actions that can't be permissibly treated near-cavalierly, where an action ϕ is treated *nearly* cavalierly if there is some possible world in which some agent is objectively morally permitted to treat or regard the doing of ϕ as if it were *nearly* morally "no big deal."

To forbid agents from treating or regarding some act as nearly morally inconsequential is, again, not to require those agents to feel or believe a particular way about that act. Furthermore, it is logically possible for some agent to violate her duty to observe the Near-Cavalier Constraint for the acts admitting of such a constraint, yet fail to do so in a blameworthy way. An agent might be unaware that there is such a constraint, and thus would not plausibly be blameworthy for failing to abide by the normative demands of the constraint. Thus, the agent might be epistemically (or subjectively) morally justified in treating or regarding ϕ as if it were morally inconsequential, but she could not be *objectively* morally justified.

As in the case of torture, I claim that there is a Near-Cavalier Constraint for the act of assaulting persons, ϕ ; thus, ϕ cannot be permissibly done in a *nearly* cavalier fashion. It would be objectively wrong for an agent to assault a person by breaking her finger in a morally near-cavalier way, say, by taking it seriously but in a borderline way: one is very near *not* taking it seriously. Because a Near-Cavalier Constraint exists for ϕ , which includes all cases of assaulting persons, it extends to cases in which assaulting persons is done under widely believed permitting conditions for self-defense.

The argument from the Cavalier Constraint proceeded as follows: There is a Cavalier Constraint for ϕ , which entails that ϕ cannot be permissibly done cavalier, and ϕ could be permissibly performed cavalierly if severe overdetermination for permissibly doing ϕ was possible. Necessarily, if the Orthodox View is true, then severe overdetermination for performing ϕ is possible. Therefore, the Orthodox View is false.

The Cavalier Constraint for assaulting persons draws a line in the sand: There is a cap on the number of permissions or extent to which those permissions can normatively affect the permissibility of an act (e.g., assaulting a person). For example, suppose as before that the permissions to protect your children, yourself, and prevent catastrophe are each sufficient to permit you to assault a person in self-defense. If those considerations permitted this, they would severely overdetermine and therefore *render possibly cavalier* assaulting that person. However, suppose that if only the first two considerations permit assault, then the assault would *not* be trivialized by severe overdetermination. In that case, the third consideration (i.e., preventing catastrophe) *would be* a permission absent the other considerations, but it doesn't function as a permission here because it can't do so.³¹

| <u>Consideration</u> | <u>Permission-Affecting Status</u> |
|------------------------------|------------------------------------|
| <i>Prevent catastrophe</i> | non-affecting |
| CAVALIER CONSTRAINT | |
| <i>Protect your children</i> | affecting |
| <i>Protect yourself</i> | affecting |

Figure 1

This is analogous to a case in which the permission to prevent catastrophe does not imply that one can permissibly commit genocide: Because genocide is absolutely wrong, it cannot be morally permitted, and thus no reason can stand in the relevant permitting relation. The Cavalier Constraint makes severe overdetermination impossible, and therefore makes impossible any and all relations that would entail such severe overdetermination.

The objector's move is to make possible *some* overdetermination with respect to assault, but not enough to render it permissible to do cavalierly. If successful, this move would for example effectively allow certain moral considerations to permit assault, but would not permit another consideration from

31. Thus, acting so as to prevent catastrophe is a pro tanto permission in the following sense: it is normative or has normative implications, other things being equal. Absent special conditions, agents are permitted to prevent catastrophe. Alternately or in addition, acting so as to prevent catastrophe is a pro tanto permission in the following sense: agents are permitted to prevent catastrophe, but not all avenues to prevent catastrophe are permissible (e.g., genocide).

simultaneously doing so. It is helpful to think of the constraints as thresholds, and Figure 1 (above) is helpful in illustrating the relations.

Above the constraint, which is rightly conceived and effectively illustrated as a threshold, no permissions normatively affect ϕ . Beneath the constraint threshold, moral considerations can normatively affect ϕ , even overdetermining the permission for doing ϕ .³² Because of the Near-Cavalier Constraint, however, our scale must include a secondary threshold, as represented in Figure 2.

| <u>Consideration</u> | <u>Permission-Affecting Status</u> |
|----------------------------------|------------------------------------|
| <i>Prevent catastrophe</i> | non-affecting |
| CAVALIER CONSTRAINT | |
| <i>Protect your children</i> | non-affecting |
| <i>Protect yourself</i> | non-affecting |
| NEAR-CAVALIER CONSTRAINT | |
| <i>Keep your promises</i> | affecting |
| <i>Oppose terrorist activity</i> | affecting |

Figure 2

Suppose now that preventing catastrophe remains above the Cavalier Constraint, which entails that it *does not affect* the permissibility of ϕ . Unfortunately for the permissions to protect your children and to protect yourself, while they did fall below the Cavalier Constraint, they would (if they affected ϕ) permit an agent in treating or regarding ϕ as *nearly* morally inconsequential. Thus, they fall above the Near-Cavalier Constraint, and therefore do not affect the permissibility of ϕ .

A proponent of the Orthodox View might and should accept this scale. No reason can affect the permissibility of ϕ if doing so would either permit ϕ to be performed cavalierly or *nearly* cavalierly. The trouble is “nearness,” which is a paradigmatic example of a predicate that necessarily admits of vague cases. Figure 2 (above) represents protecting humanity and keeping your promises as permissions normatively effect ϕ , but not in a way that would objectively permit any agent to regard doing ϕ as nearly morally inconsequential. Yet among the reasons that would permit doing ϕ (or would further permit doing ϕ), there are some considerations are metaphysically indeterminate with respect to affecting ϕ , as represented in Figure 3.

Some considerations are weighty enough that they will nearly permit ϕ to be done cavalierly, whereas others are weak enough that they will be consider-

32. Because specifying where the threshold ‘begins’ is a matter of specifying the total set of sufficient conditions for rendering an act permissible to perform cavalierly, and because this essay focuses on only one sufficient condition, no more will be said about the finer details of the constraint in regards to a scale as represented in Figure 1 (above).

ably distant from doing that. Still other duties will be borderline cases: it will be neither true nor false that they *nearly* permit ϕ to be done cavalierly. This occurs in ordinary cases of nearness. If two objects are within half an inch of each other, they are near each other. If they are a mile away from each other, they are not near each other. But it is indeterminate whether they are near each other if they are, say, forty-eight feet apart.

Because the permission to get answers from the terrorist is a similarly borderline case, it might be represented as *either* affecting *or* non-affecting with respect to ϕ . Of course, it *can't* be the case that the permission *both affects ϕ and fails to affect ϕ* , since the law of excluded middle holds for whether 'Get answers from the terrorist' is *permission-affective* with respect to ϕ . Yet as an indeterminate case, 'Get answers from the terrorist' would do just that. Such a status is logically impossible (since excluded middle holds) and, therefore, it is impossible for such indeterminate cases to exist. But as a matter of conditional necessity, there are such cases if there is a Near-Cavalier Constraint and if any reasons whatsoever normatively affect ϕ : some will affect ϕ , some will not, and still others (the borderline cases) will do neither.

| <u>Consideration</u> | <u>Permission-Affecting Status</u> |
|---------------------------------------|------------------------------------|
| <i>Protect your children</i> | non-affecting |
| <i>Protect yourself</i> | non-affecting |
| <i>Get answers from the terrorist</i> | non-affecting(?) |
| NEAR-CAVALIER CONSTRAINT | |
| <i>Get answers from the terrorist</i> | affecting(?) |
| <i>Protect humanity</i> | affecting |
| <i>Keep your promises</i> | affecting |

Figure 3

A final objection bears similarity to the last objection. There exists a view, much like the one defended here, which concedes the moral impossibility of permissibly assaulting someone in a cavalier manner. But unlike the view defended here, the alternate view denies that the Cavalier Constraint rules out overdetermination. The view is similarly motivated: permissibly performing certain actions cavalierly appears morally impossible; therefore, it is morally impossible. If self-defensive assault is permissible at all, then overdetermined self-defensive assault is possible. But this, an objector might argue, *doesn't* rule out overdetermination, even massive or infinite overdetermination. Rather, it rules out *cavalier-making* overdetermination, including cavalier-making massive and infinite overdetermination. In effect, the alternate view claims that no matter how seriously a permission to self-defensive assault is overdetermined, it remains impossible to trivialize assaulting persons.

The alternate view's deficit concerns the implausibility of maintaining that the Cavalier Constraint is necessarily unaffected by the various relations there are, including overdetermination. The truth of the Cavalier Constraint fails to rule out as impossible any actual (or merely logically possible) relations, but it rules out certain relations that, absent the constraint, would be possible. Think of this in terms of moral theory. Some moral theories do not include the Cavalier Constraint, whereas others do. Of those that do not, certain moral relations, like *permitting-cavalier-treatment-via-overdetermination*, are maintained to be possible. Thus, the Cavalier Constraint implies that some moral theories are false.

Now consider an analogy. Suppose it is claimed that the act of rotating a ceramic penguin to the slight annoyance of a friend cannot be done in a morally casual way. Against this, it is reasonably argued that such an act can surely be done with moral ease. After all, if rotating the penguin is all that is required to prevent every bad state of affairs in the world (every war, every labored and asthmatic breath), surely some agent would be objectively justified in rotating the penguin without a second thought. The evidential basis for inferring trivializing is apparent: massive overdetermination. Said another way, the worse not- ϕ -ing gets, the harder it is to justify not- ϕ -ing. Thus, the worse not- ϕ -ing, the *easier* it is to justify ϕ -ing. This is evidence for the conclusion that Cavalier Constraint is sensitive to overdetermination in the following way: Massive overdetermination is *sufficient* for permitting an action to be performed in a morally cavalier way. But that is just what the defender of the alternate view denies. Thus, the alternate view is false.

§5 Summing Up

For various reasons, most adopt and defend a particular thesis about self-defensive assault, namely:

The Orthodox View: Some instances of self-defensive assault are morally permissible.

Here it was argued that the Orthodox View is false in virtue of an untoward modal implication: that there can be moral reasons to assault a person such that those reasons permit an act of self-defense, and that the reasons which permit self-defensive assault can overdetermine the permissibility for self-defensive assault. Such overdetermination is impossible for acts for which there is a Cavalier Constraint, and there is such a constraint for assaulting persons. Since the Orthodox View implies that there is some possible world in which assaulting persons can be permissibly trivialized and there is no such world, it is necessarily false.

Against this, it was argued that the argument from trivializing itself permits advocates of the Orthodox View to reject cavalier-making overdetermination while accepting overdetermination. But this is implausible because at least for the act of assaulting persons, there is also a Near-Cavalier Constraint, and *that* constraint forces the defender of the Orthodox View to claim that there are metaphysically indeterminate cases of the following sort: some moral consideration, such as ‘Get answers from the terrorist,’ neither *normatively permits* self-defense nor *fails to permit* self-defense. Since every moral consideration necessarily does one or the other, this is impossible, and thus there can be no moral considerations that normatively permit self-defense. Plausibly, this implies that self-defensive assault is not the sort of thing that can be permitted because it is impossible, and the reason it is impossible is that it violates at least one side-constraint on permissible action, making it absolutely wrong. Thus, the Orthodox View is false.³³

Department of Philosophy, University of Washington
sbhereth@uw.edu

References

- Audi, Robert. 2009. “On the Meaning and Justification of Violence,” in *Violence: A Philosophical Anthology*, ed. Vittorio Bufacchi, 136–67. New York: Palgrave Macmillan.
- Brummer, James. 1996. “The Structure of Residual Obligations.” *Journal of Social Philosophy* 27(3): 164–80. <https://doi.org/10.1111/j.1467-9833.1996.tb00259.x>
- Bufacchi, Vittorio, and Jean Maria Arrigo. 2006. “Torture, Terrorism, and the State: A Refutation of the Ticking-Bomb Argument.” *Journal of Applied Philosophy* 23(3): 355–73. <https://doi.org/10.1111/j.1468-5930.2006.00355.x>
- Davion, Victoria. 1990. “Pacifism and Care.” *Hypatia* 5(1): 90–100. <https://doi.org/10.1111/j.1527-2001.1990.tb00392.x>
- Doggett, Tyler. 2011. “Recent Work on the Ethics of Self-Defense.” *Philosophy Compass* 6(4): 220–33. <https://doi.org/10.1111/j.1747-9991.2010.00382.x>
- Draper, Kai. 1993. “Fairness and Self-Defense.” *Social Theory and Practice* 19(1): 73–92. <https://doi.org/10.5840/soctheorpract19931913>
- Draper, Kai. 2009. “Defense.” *Philosophical Studies* 145(1): 69–88. <https://doi.org/10.1007/s11098-009-9387-5>
- Dworkin, Ronald. 1984. “Rights as Trumps,” in *Theories of Rights*, ed. J. Waldron, chap. 6. New York: Oxford University Press.
- Fiala, Andrew. 2014a. “Pacifism.” *Stanford Encyclopedia of Philosophy*.
- Fiala, Andrew. 2014b. “Pacifism and the Trolley Problem.” *The Acorn* 15(1): 33–41. <https://doi.org/10.5840/acorn20141514>

33. My thanks to Helen Frowe, Bekka Williams, and two anonymous reviewers at *Social Theory and Practice* for their comments on earlier drafts of this paper, and to Saba Bazargan-Forward for a very productive conversation on the central argument.

- Frowe, Helen. 2010. "A Practical Account of Self-Defence." *Law and Philosophy* 29: 245–72. <https://doi.org/10.1007/s10982-009-9062-1>
- Frowe, Helen. 2011. *The Ethics of War and Peace: An Introduction*. New York: Routledge.
- Frowe, Helen. 2014. *Defensive Killing*. New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199609857.001.0001>
- Kamm, F. M. 2012. *The Moral Target: Aiming at Right Conduct in War and Other Conflicts*. New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199897520.001.0001>
- Kellenberger, James. 1987. "A Defense of Pacifism." *Faith and Philosophy* 4(2): 129–48. <https://doi.org/10.5840/faithphil19874219>
- Lang, Gerald. 2014. "Why Not Forfeiture?," in *How We Fight*, ed. Helen Frowe and Gerald Lang, 38–61. New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199673438.003.0003>
- Lazar, Seth. 2009. "Responsibility, Risk, and Killing in Self-Defense." *Ethics* 119(4): 699–728. <https://doi.org/10.1086/605727>
- McMahan, Jeff. 1994a. "Innocence, Self-Defense and Killing in War." *The Journal of Political Philosophy* 2(3): 193–221. <https://doi.org/10.1111/j.1467-9760.1994.tb00021.x>
- McMahan, Jeff. 1994b. "Revising the Doctrine of Double Effect." *Journal of Applied Philosophy* 11(2): 201–12. <https://doi.org/10.1111/j.1468-5930.1994.tb00109.x>
- McMahan, Jeff. 1994c. "Self-Defense and the Problem of the Innocent Attacker." *Ethics* 104: 252–90. <https://doi.org/10.1086/293600>
- McMahan, Jeff. 2011. *Killing in War*. New York: Oxford University Press.
- Nagel, Thomas. 2005. "Ruthlessness in Public Life," in *Mortal Questions*. New York: Cambridge University Press.
- Narveson, Jan. 1965. "Pacifism: A Philosophical Analysis." *Ethics* 75(4): 259–71. <https://doi.org/10.1086/291549>
- Narveson, Jan. 2003. "Terrorism and Pacifism: Why We Should Condemn Both." *International Journal of Philosophy* 17(2): 157–72. <https://doi.org/10.5840/ijap200317217>
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
- Orend, Brian. 2005. "War." *Stanford Encyclopedia of Philosophy*.
- Otsuka, Michael. 1994. "Killing the Innocent in Self-Defense." *Philosophy & Public Affairs* 23(1): 74–94. <https://doi.org/10.1111/j.1088-4963.1994.tb00005.x>
- Parfit, Derek. 2011. *On What Matters*, vol. 1. New York: Oxford University Press.
- Quong, Jonathan. 2009. "Killing in Self-Defense." *Ethics* 119: 507–47. <https://doi.org/10.1086/597595>
- Quong, Jonathan. 2012. "Liability to Defensive Harm." *Philosophy & Public Affairs* 40(1): 45–77. <https://doi.org/10.1111/j.1088-4963.2012.01217.x>
- Quong, Jonathan, and Joanna Mary Firth. 2012. "Necessity, Moral Liability, and Defensive Harm." *Law and Philosophy* 31(6): 673–701. <https://doi.org/10.1007/s10982-012-9135-4>
- Rodin, David. 2002. *War and Self-Defense*. New York: Oxford University Press. <https://doi.org/10.1093/0199257744.001.0001>

- Ryan, Cheyney. 1983. "Self-Defense, Pacifism, and the Possibility of Killing." *Ethics* 93(3): 508–24. <https://doi.org/10.1086/292463>
- Steinhoff, Uwe. 2006. "Torture: The Case for Dirty Harry and against Alan Dershowitz." *Journal of Applied Philosophy* 23(3): 337–53. <https://doi.org/10.1111/j.1468-5930.2006.00356.x>
- Steinhoff, Uwe. 2012. "Legalizing Defensive Torture." *Public Affairs Quarterly* 26: 19–32.
- Steinhoff, Uwe. 2013. *On the Ethics of Torture*. Albany: State University of New York Press.
- Thomson, Judith Jarvis. 1991. "Self-Defense." *Philosophy and Public Affairs* 20(4): 283–310.