

HELEN FROWE

A PRACTICAL ACCOUNT OF SELF-DEFENCE

(Accepted 22 September 2009)

ABSTRACT. I argue that any successful account of permissible self-defence must be action-guiding, or *practical*. It must be able to inform people's deliberation about what they are permitted to *do* when faced with an apparent threat to their lives. I argue that this forces us to accept that a person can be permitted to use self-defence against Apparent Threats: characters whom a person reasonably, but mistakenly, believes threaten her life. I defend a hybrid account of self-defence that prioritises an agent's subjective perspective. I argue that it is sufficient to render the use of defence permissible if an agent reasonably believes that (a) she is morally innocent, and (b) if she does not kill this person, then they will kill her. I argue that the correct account of self-defence must distinguish between whether an agent is permitted to inflict harm, and whether the target is liable to bear that harm.

I. INTRODUCTION

Moral theorists are often identified as subjectivists or objectivists. What this means, roughly, is that they prioritise different sorts of facts when thinking about what morality permits us to do. Objectivists are concerned with *facts of the matter*. Some strands of consequentialism fall under this banner. A person who thinks that morality requires us to actually maximise the good will think that a given action is right if and only if it does in fact maximise the good. This will be true even if, for example, the agent had no intention of maximising the good, but intended rather to decrease overall welfare.

Other types of consequentialist emphasise the importance of expected consequences. If an agent *believes* that performing an action will maximise the good, then it is the right thing to do. This is true even if, for some unforeseen reason, the action does

not in fact maximise the good. It might have disastrous consequences that seriously decrease the good. But on such accounts, the rightness of an action is to be measured in terms of facts about what an agent believed, not facts about how the action turned out.

These two positions give us the general gist of the debate between objectivists and subjectivists. Roughly, objectivists care about how the world really is. Subjectivists care about how the world appears to agents. As far as I can tell, the importance of this divide has been rather overlooked by philosophers working in the field of self-defence. In this paper, I show that this debate throws into stark relief some of the most central problems in defensive theory. How we answer these problems will play a important role in shaping our answer to the question of what a person is permitted to do in order to save her own life.

In section one of the paper, I explain the problem of Merely Apparent Threats: characters whom Victim reasonably, but mistakenly, believes to pose a threat to his life. (From now on, I will refer to these as Apparent Threats.) Such characters are epistemically indistinguishable from Genuine Threats, who really *do* pose a threat to Victim's life. I argue that this poses a significant problem for anyone who wants their theory of self-defence to be *practical*: to be able to tell Victim what he is permitted to *do* when faced with an apparent threat to his life. Such practicality should be an aim of all theories of permissible defence. I claim that the correct account of self-defence must be hybrid: it must incorporate both objective and subjective elements. In section two, I consider Kimberly Kessler Ferzan's attempt to develop just such an account in light of Apparent Threat cases. I argue that Ferzan's account fails because it takes the objective element to precede the subjective element, and thus cannot meet the practicality requirement that I defend and that Ferzan herself endorses. I suggest that we deal with Ferzan's concern that prioritising subjectivism in self-defence will lead to general moral subjectivism by making a special case for self-defence, based upon the typically urgent nature of defensive action. In the third section of the paper, I develop and defend

a new sort of hybrid account, upon which Victim's subjective perspective suffices to justify his use of defensive force. However, I argue that different, objective factors determine whether Victim's target is liable to such force. The justness of inflicting harm is independent of the justness of the harm itself. This distinction helps us to answer not only the question of what Victim may do to an Apparent Threat, but also what that Apparent Threat may do to Victim.

II. THE PROBLEM

1. *Some Cases*

Let us begin, then, with some thought-experiments that will demonstrate the sort of problem that I have in mind. Consider *Culpable Genuine Threat*:

<i>Culpable Genuine Threat</i>	Victim is being pursued by Murderer, who has a gun and wants Victim dead because he hates Victim. Only by killing Murderer can Victim save his own life.
------------------------------------	--

I believe, along with most people, that Victim is permitted to kill Murderer to save his own life. Common explanations for this permission include the idea that there is greater utility in the survival of an innocent person than in the survival of a guilty person. Or, we might argue that Murderer has somehow forfeited his right not to be killed by maliciously attacking Victim.¹ Or, we might point out that Murderer has deliberately forced Victim to choose between their two lives and that given this, Victim is entitled to choose his own life over Murderer's.² All these explanations cite Murderer's fault, or moral responsibility for the threat to Victim's life, as crucial to the explanation of why Victim may kill Murderer.

¹ See e.g. Judith Thomson, 'Self-Defence', *Philosophy and Public Affairs* 20(4) (Autumn, 1991): 283–310, Suzanne Uniacke, *Permissible Killing* (Cambridge: CUP, 1994), and Fiona Leverick, *Killing in Self-Defence* (Oxford: OUP, 2006).

² See Cheyney Ryan, 'Self-Defense, Pacifism, and the Possibility of Killing', *Ethics* 93(3) (Jan., 1983): 508–524, and Richard Norman, *Ethics, Killing and War* (Cambridge: CUP, 1995).

With this in mind, consider *Culpable Apparent Threat*:

<i>Culpable Apparent Threat</i>	Victim is being pursued by Incompetent Murderer, who hates Victim and wants him dead. Both Victim and Incompetent Murderer mistakenly believe Incompetent Murderer's gun to be loaded.
-------------------------------------	--

I suspect that many people's intuition will be that, at the very least, Victim does not act wrongly if he shoots Incompetent Murderer on the basis of his belief that Incompetent Murderer is trying to kill him. Of course, if Victim knows that the gun is not loaded, he is not permitted to kill Incompetent Murderer. But given that Victim believes that Incompetent Murderer is going to kill him, and Incompetent Murderer certainly intends to kill him, we might think that Victim is permitted to kill Incompetent Murderer. After all, from Victim's perspective, Incompetent Murderer is indistinguishable from Murderer.

Our intuitions might be very different, though, in *Innocent Apparent Threat*:

<i>Innocent Apparent Threat</i>	Victim happens across Actor who, unknown to Victim, is rehearsing for his gangster play. Not noticing Victim, Actor points the replica gun in Victim's direction whilst saying the line, 'I'm going to kill you'.
-------------------------------------	---

In this case, Actor's innocence, along with the fact that he doesn't really pose a threat to Victim, makes it very hard to endorse the idea that Victim has some sort of moral permission to nonetheless shoot Actor. Actor bears neither moral nor causal responsibility for a threat to Victim's life, since there *is* no threat to Victim's life. But, of course, the same is true of Incompetent Murderer. Incompetent Murderer poses no more of a threat than Actor, and thus can be no more responsible, either morally or causally, for a threat to Victim's life.

The problem, then, is this. It seems that Victim should be permitted to kill at least some of the characters described above. But neither the objectivist nor subjectivist line seems to cut across the cases as we would like. The objectivist, on the one

hand, will not allow Victim to kill Incompetent Murderer, despite the fact that Victim cannot distinguish Incompetent Murderer from Murderer, and Incompetent Murderer intends to kill Victim. There is general consensus that permissible defence must be both proportionate and necessary. Defence against either Actor or Incompetent Murderer is objectively unnecessary, and thus impermissible on an objectivist account.

The subjectivist, on the other hand, will permit the killing of both Murderer and Incompetent Murderer, since from Victim's perspective both killings seem proportionate and necessary. But the subjectivist will also permit the killing of innocent, non-threatening Actor. From Victim's perspective, Actor is indistinguishable from Murderer. Subjectivism will struggle to distinguish the permissibility of killing one from the permissibility of killing the other. In other words, subjectivist accounts seem to allow too much, and objectivist accounts too little.

III. ACTION-GUIDINGNESS AND PERMISSIBILITY

Most accounts of permissible defence are not explicitly, or perhaps not consciously, objective or subjective.³ However, they are usually developed by thinking about Genuine Threat cases, defending a particular view of what Victim is permitted to do given certain stipulated conditions (specifying, for example, that Victim will definitely be killed unless he acts, or that if he acts he will definitely survive).⁴ This tendency gives such accounts an objectivist slant, with attention focused upon what it would be permissible to do given certain facts about the outcomes of various possible actions. These are indeed interesting questions. But are they the *right* questions to be asking if one wishes to develop an account of permissible defence?

³ Thomson's account is an objectivist account in the sense that she takes what is permissible to be based on objective facts. However, she does not address the question of what to do when the facts are not as Victim perceives them to be, and so does not explicitly address the debate that I am discussing here.

⁴ See e.g. Judith Thomson, 'Self-Defence', Michael Otsuka, 'Killing The Innocent In Self-Defence', *Philosophy And Public Affairs* 23(1) (Winter, 1994): 74–94, and Jonathan Quong, 'Killing in Self-Defence', *Ethics* (forthcoming).

In a recent paper, Jonathan Quong argues for a Principle of Defensive Killing that begins by stating that, “You can permissibly kill X if X will otherwise kill you.”⁵ The principle then precedes to defend certain additional conditions that must be met, once it is established that X will indeed kill you, in order for defence to be permissible. Quong stresses that his account is not meant to deal with, “cases where the agent may believe, reasonably or not, that X is going to kill them when X is not in fact going to kill them.”⁶

[T]his paper does not address the interesting issue of when, if ever, we should be *excused* (due to panic, duress, or lack of knowledge) for an impermissible instance of killing. I am only interested in the question of when, if ever, we are *morally permitted* to kill a single person in self-defence. The paper thus does not concern itself with those cases of self-defence where one or both parties have mistaken beliefs regarding matters of fact – I assume all parties always have a full and accurate knowledge of the situation.⁷

It strikes me as a mistake to build our accounts of permissible defence around knowledge that Victim cannot have. How can these accounts tell Victim what he is permitted to do when he has less than perfect knowledge? And given that, in the circumstances in which defence arise, Victim will *always* have less than perfect knowledge, how can they play any real role in telling Victim what he is permitted to do in self-defence? These cases of ‘full and accurate knowledge’ are by their very nature artificial: the assumption on which they rely will *always* be false. This is not to say that Victim will always be mistaken about the facts. Rather, it is to say that at best Victim can have a *reasonable belief* about those facts. He can predict, perhaps accurately, the way in which events will unfold. But that one’s prediction turns out to be true does not equate to the claim that one knew what was going to happen. The myth of ‘full and accurate’ knowledge that underpins existing accounts of permissible defence *is* just a myth, and can hardly be the best starting point for the correct account of permissible defence.

⁵ Quong, ‘Killing in Self-Defence’, p. 14.

⁶ Quong, ‘Killing in Self-Defence’, p. 41, n. 25.

⁷ Quong, ‘Killing in Self-Defence’, p. 2–3.

Quong relies, like most accounts, on gesturing towards excuses and blame to deal with cases in which Victim kills a person who was not, in fact, going to kill him. And I suspect that many people will think that this the best way to deal with the problem to which I have pointed. The reason why we distinguish praise and blame from right and wrong is precisely so that we can deal with these kinds of problems. Those who do what they believe to be right might in fact act wrongly, but we do not blame such agents for their actions. The agent will be excused on the basis of, for example, their reasonable ignorance.

I think that this response fails to meet what I call the *practicality requirement* that attaches to accounts of permissible defence. Telling a person what he will or will not be blamed for is not the same as telling him what him is permitted to do. Accounts that rely upon notions of blame and excuses forget that our primary concern in self-defence should lie in informing Victim's deliberation about what he is permitted to do. It is not sufficient to assure Victim that if he is not permitted, because it turns out that the gun wasn't loaded, he will nonetheless be excused. This is partly because, as David Rodin points out, it is only justifications and not excuses that ought to function in deliberation.⁸ And it is also because one is excused for actions that were impermissible. This assurance, then, amounts to nothing more than the advice that Victim acts either permissibly or impermissibly. And accounts like Quong's require knowledge that Victim cannot have in order to tell the difference between those actions that are permitted, and those which are not. If Victim is permitted to proceed *only* if X will kill him, and Victim cannot *know* whether X will kill him, it follows that Victim cannot know whether or not he is permitted to kill X. Such accounts fail, in short, to give any useful answer to the question of when one may use defensive force. It might turn out, as a matter of luck, that Quong's account will judge a given killing permissible. But such permissibility can only be

⁸ David Rodin, *War and Self-Defense* (New York: OUP, 2002), pp. 30–31.

established after the fact of defensive force. Indeed, even ‘after the fact’ verification might not be possible – how are we to know that Murderer *would have* killed Victim, if Victim kills Murderer first? This sort of *ex post* judgement cannot function in Victim’s deliberation about what he is permitted to do.

We must clarify precisely what we are after when we think about permissible defence. Do we want an account of praise, blame and excuses that tells us how *we* may treat Victim after he uses force? Or do we want an account of permissible killing that tells Victim how *he* may treat someone who (appears to) threaten his life, bearing in mind, of course, that appearances are all that Victim will have to go on? It seems to me that we ought to want the latter, and that accounts like Quong’s, strictly speaking, can only provide the former.

Thinking about this different points us towards the fact that an account of permissible defence should be action-guiding, or *practical*. By practical, I mean that an account of permissible defence must be able to tell Victim what he is permitted to do in self-defence in advance of his actually doing it. Given this requirement, we cannot deny subjectivism the central role in determining what Victim is permitted to do in self-defence. Objectivist accounts will require Victim to act differently in cases that he cannot tell apart. If it matters which situation Victim is in, and, *ex hypothesi*, Victim cannot *tell* which situation he is in, he cannot know how he is permitted to act at any given time. Accounts that yield Victim different permissions in subjectively identical situations cannot be action-guiding in the way that I have described, and must be rejected.

In what follows, I defend a hybrid approach to self-defence, that includes both subjective and objective elements, and that meets the practicality requirement. I argue that if defensive force is ever permissible, its use must be justified on the grounds of Victim’s reasonable belief that (a) if he does not kill this person, then they will kill him, and (b) that he is innocent. However, I argue that whether or not Victim is justified in inflicting harm is not the same question as whether his target is liable to bear that harm. What renders a person liable to bear

harm are objective facts about culpability. This claim has significant implications for what a person may do to avoid bearing defensive harms.

IV. HYBRID ACCOUNTS

1. *Ferzan's Account*

Kimberley Kessler Ferzan has already attempted to deal with the problem posed by Apparent Threats by developing a hybrid account of the sort that I have in mind. I think that her account ultimately fails because it prioritises objectivism over subjectivism, and thus cannot meet the practicality requirement to which Ferzan herself subscribes. However, it will be instructive to see why Ferzan favours such an account.

Ferzan begins by undermining two of the most prominent objectivist accounts of permissible defence – those of Heidi Hurd and Paul Robinson – by drawing attention to the fact that defensive action is unavoidably an act of *risk-taking*: it is always done in the face of uncertainty.⁹ Both Hurd and Robinson allow that there can be justified risk-taking, defined as risk-taking that is not wrong even if it turns out that the risk eventuates in harm.¹⁰ Driving a car within the speed limit is an instance of justified risk-taking, where ‘justified’ indicates that it is a risk that a reasonable person would take. But when it comes to self-defence, Hurd and Robinson condemn any unnecessary use of force as unjustified. If Victim kills Incompetent Murderer, he may be excused on their accounts, but he will not be justified. His failure to produce a net social benefit, by killing one when none need die, means that Victim acts wrongly.

Ferzan accuses Hurd and Robinson of, “reaching contradictory conclusions about the very same case.”¹¹ Both permit justified risk-taking. So how can they condemn Victim’s killing

⁹ Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, *Law and Philosophy* 24 (2005): 711–749.

¹⁰ Heidi Hurd, *Moral Combat* (Cambridge and New York: Cambridge University Press, 1999) and Paul Robinson, *Structure and Function in Criminal Law* (New York: Clarendon Press, 1997).

¹¹ Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, p. 721.

of Incompetent Murderer, assuming that reasonable people would act as Victim and that his risk is thus justified?

This seems to leave us with two options. We can argue that objectivism can incorporate risk-taking, and thus need not condemn Victim's reasonable but unnecessary killings. Or, we can argue that the inability to justify Victim's unnecessary killing stems from the incompatibility of risk and objectivism. Perhaps Hurd and Robinson are mistaken in thinking that they can allow risk-taking at all. Ferzan goes for this second approach, pointing out that all risks are epistemically grounded.

That is, from a God's eye point of view, an act either will or will not happen – the risk is either 0 (no) or 1 (yes). However, because we do not have this omniscient perspective, we assess risks. And the degree of risk may change depending upon the person assessing the risk because the amount of information that any actor has may differ.¹²

Ferzan also stresses, as I have done, the need for a theory to be able to guide Victim's action:

[Victim] cannot know the outcome. Rather, he must play the odds. And his action, like any epistemically limited action, must be assessed from this viewpoint... The law must not only tell people what to do when they are certain of the result, but what to do when they face risks.¹³

Telling Victim that he is permitted to kill only those people who genuinely threaten his life is of no use from either a moral or legal perspective. Or rather, such a rule is of no use to any moral or legal system that wants to be able to inform deliberation about future defence, and not merely judge past defence.

These considerations – the difficulty of incorporating risk into objectivism, and the need to be able to guide Victim's action – lead Ferzan to reject, correctly I think, any purely objectivist stance. But nor does she wish to endorse pure subjectivism. Ferzan identifies what she thinks is an equally pressing problem for the subjectivist, namely the inability to recognise that behaviour might be reasonable yet wrong. Subjectivists claim that Victim's reasonable belief suffices to justify his killing of Incompetent Murderer. But if this is true, does, for

¹² Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 722.

¹³ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 726.

example, Alice's reasonable belief that an umbrella is hers justify her taking it even when it in fact belongs to Betty? Ferzan claims that, "it seems strange to say that Alice's taking of Betty's umbrella was the 'right' or 'permissible' thing to do."¹⁴ Given this, how can we claim that Victim's reasonable belief renders his unnecessary use of defence right or permissible? Either reasonable belief is a justification of objective wrongdoing, or it's not. And unless we want to say that Alice is justified in taking Betty's umbrella, it looks like we should say it's not.

With these problems in mind, Ferzan develops a hybrid account of permissible defence. The focus of the account is what she calls the 'triggering conditions' of self-defence. These are the conditions that, if met, give rise to Victim's right to act on the basis of his prediction about the need for defensive action. Ferzan argues that the correct triggering condition is a malicious intention to harm. Whether or not the putative threat is culpable in this way is an objective matter. Once the objective triggering condition is met, however, Victim "may act based on his prediction of future harm."¹⁵ This prediction is entirely based on Victim's perspective. Ferzan argues that once the triggering conditions are met, Victim's "subjective belief in any probability [of lethal harm] is sufficient, and this belief need not be reasonable."¹⁶ Ridding ourselves of the difficulty of defining reasonable belief is, Ferzan contends, an advantage of her account.

Ferzan summarises her view thus:

Justifiable self-defense requires both (1) the existence of objective triggering conditions and (2) the prediction that harm will occur. Thus, anytime that the actor is entitled [to] make this prediction it will be a fact of the matter that an aggressor is trying to, or at least willing to, injure the defender. If the actor is incorrect in assessing the existence of the triggering conditions, his action is *not justified*. Rather, the right to act on one's prediction obtains when the triggering conditions are satisfied.¹⁷

¹⁴ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 714.

¹⁵ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 731.

¹⁶ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 739.

¹⁷ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 747.

In a footnote, Ferzan says that when an agent mistakenly believes the triggering conditions to have been met, “a reasonableness inquiry is also not required. An honest mistake should excuse a purposeful, knowing or reckless killing. Since I believe that criminal liability should not be premised on negligence, I am not embarrassed by the question of what to do with an honest but unreasonable mistake regarding triggering conditions.”¹⁸

Ferzan’s account differs from objectivist accounts because she will permit Victim to kill Incompetent Murderer, whose culpability will satisfy the triggering conditions. But she differs from the subjectivist because she will not allow the killing of innocent Actor, which will only be excused on her account. How successful is her account?

2. *Predicting Triggering Conditions*

I disagree with Ferzan on various counts. But I think that the main problem with her position is that it fails in what she herself marks out as a vital ingredient in accounts of permissible defence. The law, she says, must be able to tell Victim what he is allowed to *do*, both when he is certain of the result and when he takes a risk. But if culpability is to be assessed objectively, Victim cannot know whether the triggering conditions for his predictive right have been met. So how is he to know whether he can make his prediction? How can he know what he is permitted to do?

Victim’s belief that Murderer satisfies the triggering conditions is also a prediction: he must, if he is following Ferzan’s advice, consider the chances that the person approaching him is morally culpable with respect to the threat that she appears to pose. Ferzan says that, “self-defense is [Victim’s] right to act on his prediction and thus [Victim’s] right to sometimes be wrong.”¹⁹ But whether or not he has Ferzan’s right to make a prediction must in turn be based upon a prediction that the apparent threat is culpable. The right to make a prediction regarding defence cannot get off the ground unless Victim makes a *prior* prediction about culpability. So why doesn’t

¹⁸ Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, p. 741, n. 85.

¹⁹ Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, p. 748.

Victim's right to be wrong also cover the making of the prior prediction about culpability, which is a necessary pre-requisite of predicting the necessity of defence? And if it doesn't, how can Victim know what he is allowed to do?

I think that Ferzan is right about the problem facing objectivism. But I cannot see that she has done anything but move the bump in the carpet. The criticism she levels at objectivism is that it unfairly condemns Victim's risk taking, because "[t]he objectivist must acknowledge that the defender acts under uncertainty, yet simultaneously condemn him if he guesses wrong."²⁰ Ferzan changes which guess can be the wrong guess. Victim's justification no longer rides upon the guess about whether defence is necessary, but upon the guess about whether the target is culpable. But this guess *is also a guess*, and should Victim guess wrongly, he "*is not justified*" on Ferzan's account. Her Victim also acts under uncertainty, and is condemned by Ferzan if he guesses wrong.

3. *The Special Case of Self-Defence*

Ferzan develops her account of triggering conditions as an attempt to stave the tide of moral subjectivism. Morality must be able to allow for reasonable wrongdoing, and so total subjectivism must be dismissed. But I think that this problem of resisting general moral subjectivism can be solved by showing that we should treat defence as a special case within morality.

What morality can require of us is surely constrained by what it is possible for us to do, and these constraints are especially in force in the circumstances surrounding self-defence. Defence is by its nature *urgent*: it does not allow for the deliberation or investigation that we ought to require in other parts of morality. In Uwe Steinhoff's words, "the sudden death threat and the perceived time pressure will in many if not most cases severely disturb the ability to deliberate rationally if they do not completely undermine it":

If, especially at night in a dark alley, one has a gun pointed at one's head and hears: 'Give me you f... money or I'll f... kill you', one will hardly engage in deliberations of this kind: 'The probability that his threat is meant

²⁰ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 716.

seriously lies around x; my options to defend myself or to escape are this and that, with this and that probability of success... Rather, many people threatened with deadly violence will think: 'Oh my God, he wants to shoot me, he wants to shoot me... I'll give him the money!'²¹

I think that usually there can be *some* rational deliberation about defence. Most people who are, for example, about to fire a gun, will not do so lightly. They may look round for an escape route, or to see if they can find a place to hide. Most people will find escape preferable to killing, even the justified defensive killing of a culpable aggressor. But while they may undertake investigations about necessity regarding their environment, it will hardly be possible for them to undertake investigations about necessity regarding the putative assailant and his weaponry. What is distinctive about self-defence, as opposed to many other aspects of morality, is that demanding a high epistemic standard is simply incompatible with the necessity condition that governs defence.

Consider the two possible triggering conditions that Ferzan offers: culpability, and actually posing a danger. What determines an assailant's culpability – Ferzan's preferred triggering condition – is his mental state, his "evil intention."²² This is simply inaccessible to Victim. Even a person shouting threats or abuse might be drugged or hypnotised. How, then, can Ferzan expect her account to tell Victim what he is permitted to do, faced as he is with this apparent risk of being killed and unable to divine the culpability of his attacker? Ferzan does no better, on this score, than the purely objectivist account.

What about the second possible trigger: that threat really does pose a danger, as favoured by pure objectivist accounts? In the roulette cases under discussion in Ferzan's paper, uncovering this will involve examining the contents of the gun's chamber. But if one can examine the contents of an assailant's gun, there is clearly no need to use self-defence. Imagine a person of normal faculties who finds herself faced with what looks like a man holding a gun and advancing towards her. What can she do to

²¹ Uwe Steinhoff, 'Torture – The Case for Dirty Harry and against Alan Dershowitz', *Journal of Applied Philosophy*, 23(3) (2006): 337–353, p. 340.

²² Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 731.

determine the necessity of defensive force? Well, she can look to see if there are any possible escape routes (do objectivists condemn those who overlook even well-hidden escape routes?). She can try to check that the person really does have a gun, and not a hairdryer. But how can she ascertain this? By moving closer? By asking? By waiting until a shot is fired? Surely we don't want to demand that people wait to be shot at before using defence. Imagine a shot *is* fired, but misses. How does she know that the next one won't be a blank (or that there will be a next one at all)? Well, only by looking in the gun's chamber. But ascertaining this is simply incompatible with using self-defence. Clearly, anyone who can take their assailant's gun and examine it has no need of defensive force. Waiting until a shot is fired will often be incompatible with defence, since you can't engage in self-defence when you're dead. Waiting for the assailant to get closer might also be incompatible with defence, since it will often remove the opportunity for action.

I have suggested that it is an essential element of an account of permissible defence that it be action-guiding: able to tell Victim what he is permitted to do. And I think that in order for it to fulfil this criterion, it must be subject to a lower epistemic standard than other actions. This standard, low in comparison to those governing other actions, can be justified by the fact that it is the highest possible standard *compatible with performing defensive action*. As T. M. Scanlon argues, the idea is not that permissibility rests simply on, "what a particular agent believes the facts to be." Rather,

It depends also on what it is reasonable for the agent to believe in the situation, what it is reasonable for the agent to do to check those beliefs, and whether the agent has done those things.²³

What it is reasonable to expect a person to find out before performing an action must fall within the limits of what it is possible for her to find out compatible with performing that action. Raising the epistemic bar any higher will generate an impossibility within self-defence, since meeting those standards will preclude meeting the necessity condition.

²³ T. M. Scanlon, *Moral Dimensions* (Cambridge, MA: Belknap Press, 2008), p. 52.

The inaccessibility and incompatibility of the information that Ferzan and the objectivists respectively require forces us, I think, to accept that if we want our theory to tell us what we are allowed to do to an apparent threat in self-defence, we must look to our own liability, and not to theirs. The starting point cannot be ‘Are they liable to be killed?’. It must be ‘Am *I* liable to be killed?’. Whatever else we wish to say must follow from there.

V. INFLICTING HARM AND THE HARM ITSELF

1. *Subjective Permissibility and Objective Liability*

In order to resolve the problem of Apparent Threat cases, we must distinguish between the *justness of inflicting a harm*, and *justness of the harm itself*. This enables us to combine both subjective and objective perspectives in a single theory. We can hold that Victim’s reasonable belief that if he doesn’t kill this person, then they will kill him, coupled with his reasonable belief that he is innocent, suffices to justify the inflicting of harm: of Victim’s using defensive force. But whether or not the target is liable to be harmed – whether the ‘harm itself’ is just – will depend upon objective facts about whether the target is party to a relevant wrongdoing.²⁴ Such wrongdoing will involve a deliberate attempt to harm Victim, or a willingness to harm Victim, where this harm is not itself intended to be defensive. Of course, Victim cannot know whether or not his target is party to such a wrongdoing. But on the account that I am defending, Victim does not need to know this for the purposes of his defence.

Since Victim has a reasonable belief in all of our hypothesised cases that his life is threatened by the person he proposes to kill, Victim is justified in inflicting harm in each. If we want an account of self-defence that can guide Victim’s deliberation about future actions, and not merely our judgements about his past actions, we must accept that we cannot require Victim to act differently in cases that he cannot tell apart. If Victim is to

²⁴ The second label of my distinction – the justness of the harm itself – might raise a few metaphysical hackles. Some people think that a harm is not the sort of thing that can be just or unjust. But all that is meant by the claim that ‘the harm is unjust’, is that ‘the target is not liable to bear the harm’. Readers who object to my label can substitute this for ‘unjust harm’.

be permitted to use force in any of the cases, he must be permitted to use force in all of the cases.²⁵

However, it does not follow from the justness of Victim's inflicting harm that the target is liable to bear that harm. Mere reasonable belief on Victim's part cannot render someone else liable to bear lethal harm, especially when that belief might be mistaken. It is here that culpability matters. I suggest that it is sufficient to render a person liable to bear a lethal harm, such that their bearing the harm is not unjust, if they are culpably trying to bring about Victim's death, even if their attempt will fail. On my view, such a person is liable to be killed even if Victim *knows* that their attempt will fail (but of course, Victim would not be permitted to inflict harm upon such a person because he lacks the belief in the necessity of doing so).

On this account of liability to defensive harms, Murderer and Incompetent Murderer will not bear unjust harms. Victim does not wrong Incompetent Murderer in killing her, despite the fact that Incompetent Murderer poses no threat. Both Murderer and Incompetent Murderer are maliciously intending to kill Victim, and are thus liable to bear the harm that Victim will inflict. This is also true in cases where a person maliciously exposes Victim to a risk of lethal harm. Consider the following case:

Roulette Risky Murderer is forcing Victim to play Russian Roulette. Risky Murderer loads five bullets into his six-chambered gun in full view of Victim. He then spins the chamber and points the gun at Victim's head. Unbeknown to Risky Murderer, Victim has his own gun that he can use to kill Risky Murderer.

In *Roulette*, Victim knows that his defence may well be unnecessary. But this does not seem to affect the permissibility of killing Risky Murderer. A five-in-six chance that Risky Murderer will kill him seems sufficient to justify Victim's using force. In fact, a three-in-six chance or even a one-in-six chance seem sufficient. But what about:

²⁵ It is of course possible to insist that Victim is not allowed to use force in any of the cases I have described. I will not argue against this claim: my argument is aimed at those who take killing Murderer to be permissible.

Apparent Roulette Risky Murderer places a sixth bullet in the chamber of his gun in front of Victim. He then turns his back to Victim and empties all the bullets from the chamber. He intends only to scare Victim when he turns round and places the gun against Victim's temple.

Here, as in *Culpable Apparent Threat*, there is no real threat to Victim's life. But it is a hard pill to swallow that Victim acts wrongly – that he does something morally impermissible that wrongs Risky Murderer – if he shoots Risky Murderer. It is my contention that Victim does nothing impermissible if he kills Risky Murderer, and that he does not wrong Risky Murderer if he kills him. Risky Murderer renders himself liable to be killed in *Apparent Roulette* because he intentionally makes Victim believe that his life is at risk. And surely the same is true of Incompetent Murderer. The only difference between the two cases is that Risky Murderer knows that his threat is a bluff. But this difference cannot make a difference to what Victim is permitted to do to them. If Victim is permitted to kill Risky Murderer in *Roulette*, when he believes that there is at most a five-in-six chance that he will be shot, he is surely permitted to kill Risky Murderer when he believes that he is certain to be shot. And if he is permitted to kill Risky Murderer because he believes that Risky Murderer might kill him, he must be permitted to kill Actor if he has an equally reasonable belief that Actor will kill him.

However, Innocent Apparent Threats like Actor are not liable to bear the harm that Victim will inflict upon them. Whilst Victim's reasonable belief that his life is threatened by the person he intends to kill suffices to justify the *inflicting* of harm, this belief cannot determine whether or not someone is *liable to bear harm*. That liability results from whether or not the target culpably intends to kill Victim, or to make Victim think that they so intend. And given that Actor does not have such intent, he is not liable to bear the harm that Victim will inflict. The harm itself will be unjust, albeit justly inflicted.

This distinction amounts to the claim that one can be justified in bringing about x , even if x is not itself justified. Jeff

McMahan relies on a similar distinction in his discussion of bombing a factory surrounded by civilians. The bomber is objectively justified in destroying the factory – he is fighting a just war. But, says McMahan, “the civilians are not liable” to be killed in the attack.²⁶ The question, then, is whether the non-liable civilians may kill the justified bomber in self-defence.

The tactical bomber will... wrong them if he drops his bomb. Yet he acts with objective justification... As I noted earlier, however, it seems intuitively that the civilians are permitted to kill the tactical bomber in self-defence. How can that be if he is not liable? I suggest that people are permitted a necessary and proportionate defence against being wronged either by unjustified *or justified* action. People are not morally required, or at least not always, simply to submit to being wronged by another's morally justified action... Because the tactical bomber's justified action would wrong the civilians, they are permitted a proportionate defense. And killing him would be proportionate. But because he is not morally liable to be killed by them, they will wrong him if they kill him in self-defense. He too, therefore is permitted a proportionate defense against their permissible defensive action. He is justified in killing them in preemptive self-defense.²⁷

McMahan's resolution of this case makes sense only against an assumption that one can be justified in bringing about unjust harm. The bomber is justified in dropping the bomb, even though this will inflict unjust harm upon the civilians. But his justification for acting does not render the civilians liable to be killed. The two issues are distinct: that one is justified in inflicting harm does not entail that the target of harm is liable to bear it.

We can see this point is true more generally by thinking about a classic trolley problem. Most people think it permissible, if not required, that a bystander divert a runaway trolley away from five people even if it will then hit and kill one other person. But the fact turning the trolley and killing the one is justified does not entail that the one on the other track is somehow liable to be killed. Being innocently trapped on a trolley track does not render the one liable, yet it remains true that inflicting harm upon the one is permissible.

²⁶ Jeff McMahan, ‘The Basis of Moral Liability to Defensive Killing’, *Philosophical Issues* 15 (2005): Normativity, 399–400, p. 399.

²⁷ Jeff McMahan, ‘The Basis of Moral Liability to Defensive Killing’, pp. 399–400.

Of course, in both the trolley problem and McMahan's case of the tactical bomber, we are supposed to have objective justification for the killing of these non-liaible people. In the cases that I am discussing, however, we cannot appeal to something like utility to justify the killing of an innocent person or persons. Victim's killing of Actor will not maximise the good. But this difference does not affect the logical independence of liability from permissible infliction. If one could permissibly inflict only harm to which the bearer is liable, it would not be permissible to divert the trolley to save the five, the ten or the twenty. If it *is* permissible to divert the trolley in these cases, it follows that the justness of inflicting harm must be determined independently of the justness of the target's bearing that harm.

I suggest that the reason why we require objective justification, in the form of the greater good or something similar, in the bomber and trolley cases is that the targets of force in these cases do not meet one of the conditions that I have defended. The bomber does not believe that if he does not kill the civilians, then they will kill him. Nor does the bystander in the trolley problem believe that if she does not divert the trolley, the one will kill the five. We need objective justification to kill *these* people because we lack the reasonable belief that if we do not kill them, they will kill us (or some other innocent person). In the cases that I am interested, however, Victim *does* believe that the target of his force will kill him, unless he kills them. This, along with Victim's belief in his own innocence, is sufficient justification for the use of force.

The independence of the justness of inflicting harm from the justness of the harm itself helps elucidate the relationship between Victim and Actor. Actor, unlike Murderer or Incompetent Murderer, will be wronged if Victim kills him. But I do not think that this shows that Victim's use of force against Actor is impermissible. All it shows is that one can be wronged by a permissible action. But that one will be wronged has implications for what one can do to avoid that wrong. On my account of defensive killing, then, the difference between the culpable and the innocent emerges not in what Victim is permitted to do to each of *them*, but in what each of them is permitted to do to *Victim*.

VI. 'COUNTER-DEFENCE'

Considering what a Mistaken Threat is permitted to do to Victim tells us a lot about the moral status of Apparent Threats. Consider *Defensive Murderer*:

<i>Defensive Murderer</i>	Incompetent Murderer realises that her gun is unloaded just as Victim opens fire. However, another gun is lying on the ground. Only by using this loaded gun to kill Victim can Incompetent Murderer save herself.
---------------------------	--

Incompetent Murderer realises that, actually, she never posed any threat to Victim. Is Incompetent Murderer now the victim of an unjust attack, permitted to use defensive force against Victim? I think we must say that she is not. It would be an unwelcome state of affairs if merely pretending to kill a person might afford one the right to *actually* kill that person on grounds of self-defence.

If Incompetent Murderer is not permitted to defend herself, this must be because she will not suffer an unjust harm if Victim succeeds. And if the harm will not be unjust, Incompetent Murderer must be liable to bear it. As McMahan says, “[t]here can be no permissible defence against harm to which one is morally liable.”²⁸ Yet, since Incompetent Murderer did not pose a threat to Victim, it cannot be either moral or causal responsibility for a threat that has rendered her so liable. Rather, what has rendered Incompetent Murderer liable is her culpability: her malicious intention to inflict harm upon Victim. This culpability must be sufficient to render Incompetent Murderer liable to be killed if I am right that she is not permitted to fight back.²⁹ Similarly, it seems clear that Risky

²⁸ Jeff McMahan, ‘The Basis of Moral Liability to Defensive Killing’, p. 400.

²⁹ Note that this does not entail that *mere* culpability would justify Victim’s killing of Murderer. Very Incompetent Murderer, whose laughably inept plan to kill Victim poses no threat at all, is liable to be killed by virtue of his malicious intention. But it is not permissible to kill Very Incompetent Murderer in the absence of a reasonable belief in the necessity of doing so. As I have emphasised, whether someone is liable to be killed, and whether killing them is permissible, are distinct questions.

Murderer would not be able to defend himself against Victim, even in *Secret Roulette* where Risky Murderer poses no threat at all. As Ferzan says, “the person who culpably initiates the situation can hardly be heard to complain that the other actor takes her at her word.”³⁰

But despite this, I don’t think that Ferzan fully considers the implications of counter-defence for her account. We can see this by looking at what Ferzan says about Innocent Genuine Threats:

Innocent Genuine Threat Friend points a gun at Victim. Both Friend and Victim mistakenly believe the gun to be unloaded. Friend begins to pull the trigger.

Friend innocently, but genuinely, threatens Victim’s life. Ferzan argues that it is impermissible for Victim to kill Friend (even if Victim were to “come to believe” that the gun is loaded), because Friend does not meet the triggering conditions of permissible defence.³¹

In my view, one may not rightly kill these agents because they did not intentionally create the situation that forced the choice by the defender. The defender and the aggressor are both innocents. Thus, while a defender may be excused if he kills an innocent aggressor, he is not justified... Given that innocent actors do not culpably create the situation, we should not view killing them as the right thing to do.³²

I outline what I think we ought to say about innocent threats like Friend below.³³ What should concern us here is that Ferzan, despite looking at the various possible ways in which an innocent person might (seem to) threaten Victim, does not tell us what an innocent person is to do when Victim threatens *him*. Imagine that Actor suddenly notices Victim, who is about to open fire upon him. Actor has time to grab a real gun in order to defend himself. What should Actor do, on Ferzan’s account, when faced with a person who believes that he is

³⁰ Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, p. 731.

³¹ Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, p. 748.

³² Kimberly Kessler Ferzan, ‘Justifying Self-Defense’, p. 733–734.

³³ I also discuss the permissibility of killing innocent threats in Frowe, H., ‘Equating Innocent Threats and Bystanders’, *Journal of Applied Philosophy* 25(4) pp. 277–290.

engaging in permissible self-defence? Ferzan does say that the harming of either of two innocents is a tragedy.³⁴ We might infer from this, and from the passage quoted above, that she thinks, for example, that the killing of either Victim or Actor can at best be excused. But this seems inconsistent with her view about when Victim may defend himself.

Ferzan pronounces Victim's killing of a character like Actor unjustified, because Actor, being innocent, fails to satisfy the triggering conditions. But Victim is certainly not *culpably* creating a situation that forces a choice by Actor. He intentionally creates the situation, but only in the sense that his actions proceeded from his agency (just like Actor's decision to rehearse the play proceeded from *his* agency). And this isn't what Ferzan means by 'intentionally create': she has in mind *evil* intention. Victim's intention to harm, being a defensive intention, is not an evil one. Besides, it cannot be Victim's *intention* to create a forced choice if Victim's belief is that the choice is being forced upon *him* by Actor.

The problem, then, is that Victim *also* fails to satisfy Ferzan's triggering condition of culpability. On Ferzan's account, Victim is an innocent aggressor whom Actor cannot justly not kill. But Ferzan also classifies the threat that Victim poses to Actor as unjust. Victim will act wrongly in killing Actor, because Actor is innocent. This is why killing Actor is only excused. The pertinent question is whether Actor is permitted to defend himself against what will be, by Ferzan's own lights, an *unjust killing*. The intuitive answer is that he is so permitted. But I do not think that Ferzan can either grant or withhold such a permission. If she concedes Actor a permission to kill Victim, she undermines her own triggering conditions for permissible defensive killing. Victim is innocent, and thus Actor cannot have the right to make a prediction about the need for defence against him. But if Ferzan denies that Actor may prevent the inflicting of this unjust harm, this sits uncomfortably with the idea that it will *be* an unjust harm at all. Saying that Actor may do nothing to prevent Victim from killing him is very much at odds with the claim that Victim's killing of Actor is unjustified. Whilst it is certainly true

³⁴ Kimberly Kessler Ferzan, 'Justifying Self-Defense', p. 731.

that Risky Murderer may not complain if Victim tries to kill him, it is hard to see how Actor could be in the same boat. Thinking about counter-defence gives us additional reason to reject any account that insists upon a target's culpability as a pre-requisite of a permission of defence.

On my account, the fact that Actor does not culpably intend to harm Victim means that Actor is not liable to be killed by Victim. Any harm Actor suffers will be unjust. This is true even though Victim is justified in inflicting such harms. I think our intuitions strongly suggest that, faced with Victim, Actor is justified in trying to kill Victim. Just like Victim, Actor reasonably believes that he is innocent, and he reasonably believes that if he does not kill Victim, then Victim will kill him. Actor is therefore permitted to use lethal force against Victim.

VII. CONCLUSION

The aim of this paper has been to recognise an important aspect of self-defence – that defensive actions are undertaken under conditions of particular epistemic limitations, and must typically be performed urgently. Beginning with this recognition, I developed an account of permissible defence that is not only compatible with these constraints, but is in part influenced by them. I have argued that a primary concern in developing accounts of self-defence should be that it can play a practical role in a person's deliberation about what they are permitted to do. This role is importantly different from that of telling us whether a person is to be blamed or excused for a past action. If we accept the need for this practical role, we are forced to prioritise the agent's beliefs about the necessity of defence over objective facts about that necessity. I defended the priority of the subjective element of defence against Kimberly Kessler Ferzan's hybrid account, which takes objective triggering conditions to precede the right to make a prediction about the need for defence. I showed that Ferzan's account only moves the bump in the objectivist carpet, and cannot be any more practical than a purely objectivist account of permissible defence. I then developed an account of self-defence that distinguishes the permissibility of inflicting harm from the target's liability to

that harm. I argued that it is sufficient to warrant the infliction of harm if a person reasonably believes that (a) she is innocent, and (b) if she does not kill this person, then they will kill her. However, whether or not the target is liable to be killed depends on objective facts about their culpability. If they have a malicious intention to seriously harm Victim, to expose him to a risk of serious harm, or to make Victim believe that they will seriously harm him, they are liable to be killed by Victim. I suggested that liability is important for determining what a person may do to defend themselves against Victim. A person who is liable to be killed by Victim cannot use force to avoid bearing the harm that Victim will inflict. But a person who is not so liable – who lacks a malicious intention regarding Victim – is permitted to use counter-defence against Victim.

ACKNOWLEDGEMENTS

I am very grateful to The Leverhulme Trust for financial support received during the writing of this paper. I am also grateful to Jeff McMahan and two anonymous referees for extremely helpful written comments, and to Chris Hookway, Gerald Lang, Jimmy Lenman, Andréas Lind and John Oberdiek for helpful discussions. Thanks also to audiences at research seminars at the University of Kent and the University of Sheffield, and to participants at a conference on Moral Theory at Lund University.

VIII. APPENDIX

A. Reasonable Belief

My account makes use of that rather mysterious creature we call the *reasonable belief*. It is beyond the remit of this paper to develop an account of what constitutes a reasonable belief in general, beyond noting the importance of the particular epistemic conditions of defence that I have discussed above. But I ought to say something about my caveat that Victim must have a reasonable belief in his own innocence.

Let us begin by noting that on my account, actually posing a threat does not entail that one is not morally innocent.

A person might be a genuine threat to Victim's life, but the fact that she is a threat might not proceed from her agency, or indicate that she is at fault. This is the difference between *material* innocence, where one does not pose a threat, and *moral* innocence, where one is in no way culpable for a threat.

So, we can imagine a case in which Victim might unwittingly threaten another person's life. Jeff McMahan considers a case in which a terrorist has rigged Victim's cellphone in such a way that when Victim presses 'send', he will detonate a bomb that kills some innocent person.³⁵ Victim is a material threat: his actions, unbeknown to him, endanger another person's life. But Victim is not culpable with respect to the threat that he poses. Such a case illustrates the possibility that one might have a reasonable, but mistaken, belief about one's material innocence. However, the sort of reasonable belief that Victim needs on my account is a belief in his moral innocence. And it is much harder to imagine how one might be reasonably mistaken about this sort of belief.

On my account, culpability stems from intentionally exposing someone to a significant risk of serious harm that is not intended to be defensive, or that lacks 'objective' justification of the sort that we might have in the trolley problem. Less technically, culpability is going to involve intentionally, maliciously trying to harm someone, or being reckless with respect to the fact that you might harm them. It is difficult to see how one could do this and yet still reasonably believe that one is morally innocent. A reasonable belief that one is morally innocent, as I understand moral innocence here, will therefore almost always be a true belief.

B. Innocent Genuine Threats

Notable by its absence on the account that I have defended is any requirement that Victim have a reasonable belief that his target is culpable with respect to a threat to his life. All that I have committed myself to is the claim that Victim must believe

³⁵ Jeff McMahan, 'The Basis of Moral Liability to Defensive Killing', p. 397.

himself innocent. Whether or not one may intentionally kill the innocent is, I think, one of the most important questions in the ethics of self-defence, and it may be helpful to indicate my position on this issue, although it is beyond the remit of this paper to defend that position here. I think that innocent threats can be permissibly killed, and that they share a moral status with Innocence Apparent Threats like Actor: Victim may kill them, and they may kill Victim. Their innocence, along with their reasonable belief that if they do not kill Victim, Victim will kill them, suffices to justify their use of defensive force.

C. Third Party Intervention

What ought we to say about third party intervention in Apparent Threat cases? I do not have the space to fully explore this very interesting topic here. However, again, it may well be worth indicating my position. So, consider the following case, *Realisation*:

Realisation From a distance, Third Party sees Incompetent Murderer chasing Victim. Third Party knows that Incompetent Murderer's gun is unloaded. She sees Victim about to shoot Third Party, and realises that Victim thinks Incompetent Murderer is going to kill him. She can stop Victim from shooting Incompetent Murderer only by shooting Victim.

I think it impermissible for Third Party to kill Victim to prevent his killing Incompetent Murderer. Victim is not liable to be killed. But, as I have argued, Incompetent Murderer *is* so liable, and thus it cannot be permissible for Third Party to shoot Victim to save Incompetent Murderer. However, there may be a consequentialist justification for causing Victim some minor harm to prevent his shooting Incompetent Murderer.

More difficult is a case in which Third Party sees Victim about to shoot Actor. Imagine that Third Party realises that not only does Actor pose no threat, but that Actor is innocently rehearsing a play and has no culpable intent towards Victim. However, my intuition in this case is still that Third Party may not intervene, killing Victim to save Actor. Perhaps this is at

odds with what most people think it permissible to do in other-defence. But whereas self-defence can be about asking oneself whether one is liable to be killed, other-defence must work on an assumption about whether someone *else* is liable to be killed (since it is, by assumption, a case in which the third party is not herself about to suffer harm). I suspect, therefore, that there is a permission to use other-defence only when one believes that the target of the force is liable to be harmed. Since in the case we are considering, Third Party realises that both Actor and Victim are morally innocent, she cannot think either liable to be killed. She is thus not permitted to intervene on behalf of either. I think the same applies in a conflict between Victim and a Genuine Innocent Threat like Friend. When the Genuine Threat is an *innocent* threat who is not liable to be killed, a third party may not kill the innocent threat to save Victim.

Department of Philosophy
University of Sheffield,
45 Victoria Street,
Sheffield S3 7QB, UK
E-mail: H.Frowe@shef.ac.uk