https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

UNCERTAIN RIGHTS AGAINST DEFENSE

By Bas van der Vossen

Abstract: *In this essay, I defend a theory of liability to defensive force. The theory contains two elements. The first is a dual Lockean-inspired condition. The second aims to make this first condition consistent with problems arising from uncertainty. Drawing on recent work by Michael Zimmerman, I argue that the rights-based condition should be made sensitive to the evidence available to defenders.*

Suppose someone pulls out a gun and takes aim at you with a crazed look in his eyes. Having watched many Westerns, you quickly draw your own gun and kill the person in self-defense. Were you justified in doing so?

It seems we need to know more of the facts to answer this question. If you were on a movie set, and you knew that the person aiming at you was an actor, then you were not justified in killing him. If you knew that he was a dangerous murderer, and the only way to defend yourself was to shoot this person, then you would be justified. But what should we think if you were not sure about what was going on? What if the person might be an actor, but you have also been warned about a dangerous killer who is on the loose? Would you be justified in killing this person? This is the issue I want to address.

More precisely, I want to address the question of *liability* to defensive force under uncertainty.[1] The conditions of X's liability, as I understand it, are the conditions under which X forfeits certain rights, in this case rights against defensive force. I will defend what can be broadly labeled an *evidence-sensitive, rights-based* view about liability. On such a view, people become liable to defensive force when and because others have, subject to certain conditions, compelling evidence that those people threaten the rights of others, in some way to be specified.

I. The Question

Consider again the person drawing a gun on you. I said that one of the things we needed to know was whether this person is an actor or murderer. The reason is obvious. A key part of the justification of defensive force is whether the target of such force has a right that we not use it. People can lack this right because they forfeited it. The conditions under which a person forfeits rights against defensive force are the conditions under which

[1] The arguments below deal with uncertainty about the situations in which we find ourselves, not uncertainty about what is the correct moral theory.

doi:10.1017/S0265052516000133
© 2016 Social Philosophy & Policy Foundation. Printed in the USA.

he is liable to the use of such force. If someone is about to murder you, he might thereby forfeit his rights against, and thus become liable to, your defensive force. Different theories offer different explanations for why the murderer might be liable. Perhaps liability tracks culpability, moral responsibility for an unjust threat, or something else. (More on this below.)[2]

However, in order for us to be morally justified in treating the person drawing the gun as liable to our defensive force, must we *know* that he is a murderer? It is attractive to say yes, since doing so identifies as liable all and only actual murderers (and relevantly similar people). Indeed, this is the answer most philosophical theories of liability propose in one way or another. Let us call such theories of liability *objectivist* theories. On an objectivist view, the conditions of liability refer to the objective facts about the purportedly liable person(s), the defender(s), or the relation in which they stand.

Cases of uncertainty are problematic, however. When the thing we have good reason to believe to be the case may or may not in fact be the case, questions of liability become very difficult to handle in objectivist ways. If you were told by a very reliable source that a dangerous killer was on the loose, and someone pulls a gun on you, you may have good reason to believe you are in mortal danger, even if the person turns out to be only an actor. If you use defensive force in such a situation, would you violate the person's rights? This is our question.

There are four possible scenarios here. First, the person drawing a gun might actually be a killer, as indeed your evidence suggests he is. Second, the person might be an actor, and again your evidence might fit the facts. Third, the person might be an actor, while your evidence is that he is actually a killer. And fourth, he might actually be a killer, although your evidence points to him being an actor.

The first two cases are fairly straightforward. The killer is, while the actor is not, liable to defensive force. Thus, other things equal, you were justified in shooting the killer and not the actor.[3] But what about the third and fourth cases? On an objectivist view of liability, these cases, too, seem straightforward. In the third case, the person (being just an actor) is not liable to defensive force, even though you had good evidence that he was a killer. If you were to shoot the actor, you would violate his rights. You would do something morally wrong, even if your actions were understandable. In the standard language, you violated his rights, and thus acted wrongly, but may be *excused* for doing so.[4]

[2] Throughout, I will assume that the force used is proportionate. Thus, the liability we discuss is only liability to proportionate defensive force. I set aside other issues, such as whether the force is necessary, or what it exactly entails.

[3] Throughout, when discussing moral permissibility, "ought," and other similar terms, I will insert the phrase "other things equal" to indicate that I am talking about these terms to the extent that they track the presence or absence of rights.

[4] See e.g., Heidi Hurd, *Moral Combat* (Cambridge: Cambridge University Press, 1999), 7.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

The inverse applies to the fourth case. You would not act wrongly if you shoot the person because he was, in fact, a killer. This is so even though your evidence suggested he was innocent. Using objectivist language, shooting the killer would have been terribly irresponsible, and you would be *culpable* for choosing to do so. But, other things equal, you did nothing that is morally wrong. After all, the killer, being in fact a killer, was liable to defensive force.

I no longer believe that the objectivist way of dealing with these cases is satisfactory. And since in real life we are *always* less than certain about the facts, I no longer believe that the objectivist way of dealing with defensive liability in general is satisfactory. We must rely on imperfect information about other people's plans, the consequences of their actions and ours, and so on. Cases in which we are not sure in which of these scenarios we find ourselves pose the central question about the justifiability of defensive force.

The remainder of this essay consists of three parts. First, drawing on recent work by Michael Zimmerman, I will argue that objectivist approaches to liability fail and offer a basic outline of what, as far as I can tell, is the best response to this problem.[5] Second, I will discuss two ways in which others have tried to deal with problems of uncertainty and show why these fail. This reinforces the arguments from the first part. In the third part, I begin to fill out the outline of the first part with a Lockean rights-based account of liability and respond to some objections.

II. Attack under Uncertainty

Consider the following, well-known example.

> *Mistaken Attacker:* Vince's car breaks down and he knocks on Annie's door for help. When Annie answers the door and sees Vince, she believes he is about to attack her. Annie has very good evidence for this because she received a warning by the police that a villainous

[5] See Michael Zimmerman, *Living with Uncertainty* (Cambridge: Cambridge University Press, 2008) and *Ignorance and Moral Obligation* (New York: Oxford University Press, 2014). Zimmerman discusses self-defense in *Living with Uncertainty*, 108ff. There are two differences between Zimmerman's discussion and my own. First, Zimmerman offers no substantive theory of liability, but uses the case of self-defense to test his broader view. The final, rights-based part of the argument below thus serves to complement his discussion. Second, I disagree with Zimmerman on two counts. I find his arguments against objectivist theories of liability unsatisfactory because he claims that such theories cannot explain why aggressors become liable before the actual attack has occurred. This seems false. Moreover, Zimmerman says that if A threatens V, and A knows this, but V does not know this, then A nevertheless becomes liable to V's defensive force. (See Zimmerman, *Living with Uncertainty*, 108–10.) My account does not support this conclusion, as, it seems to me, is the view's natural implication. Zimmerman claims (p. 109) that V would still be wrong to attack A, but for some *other* reason than A's rights. This strikes me as ad hoc. What is wrong about V's attack if not how it treats A?

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

> murderer, whose description Vince fits perfectly, has escaped a nearby prison. Unfortunately, the evidence is wrong. Vince, by tragic coincidence, is the murderer's twin brother. Upon seeing Vince, and acting on her evidence, Annie attacks Vince in (mistaken) self-defense.[6]

The objectivist holds that since Vince was not actually threatening Annie in any significant way, Vince was not liable to Annie's defensive force. And so, other things equal, according to the objectivist, Annie ought not to attack Vince. Perhaps Annie can be excused if she nonetheless attacks Vince because her evidence conspired against her, but she did wrong Vince. She violated his rights. We can see this, the objectivist might push on, because were Annie to find out that Vince was innocent, she too would recognize that she ought to respect Vince's objectivist rights.

The problem with this way of treating *Mistaken Attacker* is what it implies about how conscientious moral agents should act. In general, it is not true that what we ought to do is what the objectivist singles out. Consider another well-known type of case (which traces back at least to Frank Jackson, and is discussed at length by Zimmerman).

> *Uncertain Rescue*: Annie has evidence that Vince is about to detonate a bomb, which will kill at least a hundred innocent people. There are three ways in which she might try to stop him. Option A is to kill Vince. Option B is to kill Vince and five of the innocents. And option C is to kill Vince and ninety-five of the innocents. Unfortunately, Annie's evidence is incomplete. It tells her only (i) how to take option B, (ii) that the remaining two options are A and C, but not which is which, and (iii) that there is no time to wait and collect more evidence.

If Annie is to act on objectivist grounds, then she ought to choose option A. The objectivist, again, might say that Annie should be excused for not doing this, given her evidence, but A remains the right thing to do. After all, as far as the facts go, options B and C involve killing innocent (nonliable) people. If what we ought to do is, other things equal, a function of the objectivist facts, then Annie ought to do A.

But this is plainly unacceptable. The only way Annie can take option A is to take a 50/50 gamble between A and C in *Uncertain Rescue*. And that is something she clearly ought *not* to do. Annie cannot permissibly take the chance of killing ninety-five nonliable people, given her other options. To take that gamble would be horribly irresponsible, even though it is the only way for Annie to do what is demanded in objectivist terms.

[6] The example is a variation on a case presented by Jeff McMahan, "The Basis of Moral Liability To Defensive Killing," *Philosophical Issues* 15 (2005): 386–405, at 387. Jonathan Quong, "Liability to Defensive Harm," *Philosophy and Public Affairs* 40 (2012): 45–77 also discusses a variation of this example (p. 53).

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

Three things are worth noting here. First, the judgment above remains true even if Annie took the gamble and things worked out favorably — by option A materializing. For Annie would still not have done what she ought to have done in light of the (non-)liability of the people involved.

Second, the judgment also remains true if we look at things from the other side. If Annie chooses B, the five innocent people who will die as a result surely have reason to regret Annie's choice. But it is implausible to say they could legitimately demand that she choose something else. Surely they, too, will recognize that, in light of her evidence, Annie made the choice she ought to have made. The same does not hold if Annie takes the gamble. The people who would have definitely been saved had she taken option B *do* have standing to insist that she act otherwise. After all, Annie put their lives in unnecessary peril.

Third, the normal objectivist responses here do not work. It is clearly mistaken to say that Annie ought to choose option A. And it is clearly mistaken to say that, were Annie to find out afterward that option A would have killed Vince alone, she would recognize that she should have taken option A. It remains wrong for her to have chosen A. Finally, it will not do to say that Annie ought to do B because she can be excused given her evidence. Excuses are not among the appropriate reasons or input for the question of what to do. What we need are justifications. If Annie can avoid doing something that she ought not to do, then she ought quite simply to avoid it. And this is true even if she would have an excuse for doing it anyway.[7]

The upshot, then, is that Annie ought to choose B in *Uncertain Rescue*. But she knows that B is not what is recommend by objectivist theories of what she ought to do. After all, she knows that B means killing five non-liable people who she could save. It is false, therefore, that we generally ought to do what objectivist theories require.

What underlies this, I believe, is that when we take up the perspective of the moral agent as such we are asking, in a first-personal manner, what she, morally speaking, ought to do. To say we take up this question in a first-personal manner is simply to say that we want to know what *we* ought to do in situations like this, acting in a morally conscientious way. We want to know how we should act in order to avoid violating other people's rights, to comply with our duties, and so on.

This may sound like a platitude. And in a sense it is. Ethics is a practical discipline; what else could it be for? But if it is a platitude, *Uncertain Rescue* shows that objectivist approaches cannot live up to it. For it is simply false to say these (platitudinous) things about the claim that it is *wrong* for Annie, given her choice-situation, to choose B. No morally conscientious person would recognize that as the correct verdict.

[7] Compare Rodin's discussion about justifications and excuses in David Rodin, *War and Self-Defense* (New York: Oxford University Press, 2002), 30–31.

Instead, what Annie ought to do is what we might call her best bet, in light of the evidence. This notion of a best bet is different from doing what is the best thing in objectivist terms, but rather refers to what is the right thing to do in terms of the evidence available.[8] Sometimes it is clear what our best bet consists in. In *Uncertain Rescue*, in light of the evidence available to Annie, and given the appropriate weighing of the liable and non-liable lives involved, Annie's best bet is option B, even though it will involve avoidably killing five non-liable people.

But if this is the right way to think about *Uncertain Rescue*, then it is also the right way to think about *Mistaken Attacker*. And there, Annie's evidence was that she was dealing with a villainous murderer. In light of this, and given an appropriate weighing of the issues concerning liability, it again seems that Annie's best bet is that she is not prohibited from defending herself against Vince. And if that was sufficient in *Uncertain Rescue* to conclude that Annie ought to choose B, then it is also sufficient in *Mistaken Attacker* to conclude that Annie can permissibly attack Vince.

We can now see the trouble for objectivist theories of liability. The conclusion that Annie is permitted to attack Vince in self-defense in *Mistaken Attacker* is precisely contrary to the verdict of objectivist theories. For if it is false that Annie ought not to attack Vince, then it is likely also false that doing so would wrong Vince. After all, it is difficult to see how, other things being equal, Annie would wrong Vince without doing something that she ought not to do. But if Annie does not wrong Vince, then she will also likely not violate a duty owed to Vince. After all, it is difficult to see how, other things equal, Annie might violate a duty owed to Vince without thereby wronging him. And if Annie does not violate a duty she owes Vince, then she likely also does not violate Vince's rights. After all, it is difficult to see how, other things equal, Annie might violate Vince's rights without violating a duty she owes him. Vince, in other words, is liable to Annie's use of force in *Mistaken Attacker*. The objectivist theory of liability is wrong.[9]

The price of denying this, again, is the unacceptable verdict in *Uncertain Rescue* that Annie ought to choose option A. So we can conclude that whether Annie's actions are wrong, or whether they wrong Vince, or violate his rights, does not depend on objectivist facts about Vince, his actions, and related issues, but on Annie's evidence concerning Vince, his actions, and related issues. In short, Vince's liability to Annie's defensive force depends on whether it is Annie's best bet in light of the evidence that Vince is liable.

Two things are worth noting about this notion of a best bet. First, the observation that objectivist theories should be replaced with theories

[8] I draw the notion of a "best bet" from Zimmerman. For more detailed discussion, see *Living with Uncertainty* and *Ignorance and Moral Obligation*, throughout.

[9] Throughout I assume that it is not Annie's fault that the evidence is faulty.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

organized around the idea of an evidence-based best bet is part of a more general thesis about what we ought to do morally. As such, the implications for the question of liability will be sensitive to the more general truth about what will count as a best bet. An attempt at providing a full account of this general moral truth would take us well beyond the confines of this essay. Instead, my aim here is to offer an account of what would go into the idea of a best bet insofar as questions of liability are concerned. Section IV below outlines such an account.

Second, the idea is not plausibly understood in terms of maximizing expected moral value, such that what we ought to do is whatever option offers the highest product of moral value and its probability of coming about. Suppose, for example, that Annie knows that there is a 90 percent chance that Vince is liable to the use of lethal defensive force. And suppose that the expected moral value of this is sufficient for Annie to be (other things equal) justified in killing Vince. If we label the net value of saving the innocents as V, then it follows that the expected value of Annie's killing Vince is 0.9*V. But compare this to a group of a hundred people, ninety of whom Annie knows with certainty are about to kill five innocent people, and thus are liable, while the remaining ten are innocent, and thus are non-liable. My intuition is that Annie would not be (other things equal) justified in killing one person at random from this group. But the expected value is the same as above: 0.9*V. So a simple expected value calculation cannot capture the intuitive notion of the morally best bet.[10]

III. Against Hybrid Theories

The argument above shows that Vince's liability is sensitive to Annie's evidence. The combination of (a) the observation that Annie did not do something she ought not to do by using force against Vince, and (b) the claim that if Vince were not liable, then (other things equal) Annie would have done something wrong by attacking him, implies that (c) Vince was (other things equal) liable to Annie's defensive force in *Mistaken Attacker*.

The two key claims in this argument can be summarized as follows. The first concerns what may be called the first-personal nature of moral requirements, "oughts," duties, and so on. This claim holds, quite simply, that the moral verdict of what Annie ought to do in situations like the cases above is what we would want *her* to *do* in those cases. We can ask whether a theory of liability delivers acceptable first-personal moral judgments by asking whether a well-informed and morally conscientious person would recognize the thing Annie is permitted or required to do as (other things equal) indeed the right thing to do.

[10] Similarly, X's best bet about Y's liability cannot be simply understood as X being justified in believing that Y is liable. While the account proposed below refers to some of X's justified beliefs, it does not follow that X's best bet can be entirely understood in this way. See note 29 below.

This claim is very plausible. We are asking when people are morally permitted to *use* defensive force, insofar as other people's rights are concerned. And to be permitted to use defensive force in this sense, it must be true that the morally conscientious person would accept that using defensive force is indeed the right or morally acceptable thing to do. The question we are asking is what the moral agent should do.[11]

The second claim concerns the connection between Vince's rights and what Annie is (other things equal) morally required to do. I take it that all theories of liability to defensive force accept something like the claim that if Annie uses defensive force against Vince, and Vince is not liable to such defensive force, then (other things being equal) Annie thereby wrongs Vince, and is therefore doing something that (other things equal) she morally ought not to do.

This claim, too, is very plausible. It expresses the correlativity of claim-rights and duties, a centerpiece of moral and legal philosophy. Indeed, it is central to the very motivation to think about liability. A primary reason we care about questions about liability is that people's rights imply duties and thus have important implications for what we ought to do. The fact that we normally have rights against force is central to the question of when defensive force is morally permissible.

The conjunction of these claims implies that Vince's rights are sensitive to Annie's evidence. This, of course, is likely to cause (considerable) discomfort. In this section I consider two theories that propose hybrid ways of dealing with the problem of uncertainty, combining objectivist and evidence-based considerations in order to avoid the discomfort. Unfortunately, these theories fail as neither can adequately incorporate both of the plausible claims above.

First, Jonathan Quong argues that people like Vince become liable to defensive harm when they do something that results in a threat of harm to people like Annie, provided she has a right against such harms, and subject to constraints about the evidence available to her. Quong distinguishes between *fact-relative* and *evidence-relative* permissibility. Acts are permissible in the fact-relative sense when they are, in fact, permissible. Fact-relative permissibility in no way depends on one's evidence or epistemic situation, but only on how things objectively are. Evidence-relative permissibility refers to an agent's evidence about whether an act is, in fact, morally permissible. When we act under uncertainty, we may not know what would be the objectively right thing to do.

[11] For related points, see Helen Frowe, "A Practical Account of Self-Defence," *Law and Philosophy* 29 (2010): 245–72; Kimberly Kessler Ferzan, "Justifying Self-Defense," *Law and Philosophy* 24 (2005): 711–49; and Kimberly Kessler Ferzan, "Culpable Aggression: The Basis for Moral Liability and Defensive Killing," *Ohio State Journal of Criminal Law* 9 (2012): 669–97.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

Quong proposes the following hybrid theory of liability:

> A person is liable to defensive harm for choosing to do X when that choice results in a threat of harm to innocent people if and only if: (a) choosing X meets the minimum conditions of moral responsibility, and (b) the evidence-relative permissibility of choosing X depends either on the assumption that those who are harmed (or might foreseeably be harmed) by choosing X are liable to the harm, or else on false moral beliefs.[12]

The main idea is that we become liable for choices that impose threats on others when those choices do not treat them as their moral status demands.

There is, however, an important instability in this theory, at least insofar as it is meant to deal with the problem of uncertainty. Quong's account suggests that Annie is liable to Vince's counter-defense in *Mistaken Attacker* (a conclusion I share, and discuss further below). However, the need to recognize the importance of evidence militates against the motivation for this conclusion. Quong says that Annie becomes liable because, even though she was permitted to attack Vince in the evidence-relative sense, she was not permitted in the fact-relative sense. And people become liable if they choose to do something that results in a threat that is, in fact, unjustified. We should hold Annie liable on fact-relative grounds, Quong writes, because she took a risk when she acted — the risk that she was treating Vince in a way that would be justified only if he lacked a right against her use of force. This makes it fair to impose the burden of the mistake on Annie.[13]

But this is unacceptable. Annie acted in accordance with her best bet in light of the evidence (and it is not her fault that the evidence is faulty). This means, we saw, that it is untenable to say that she acted wrongly. Again, what else would we want her to do? Moreover, she did not contribute more to the tragic situation than Vince did. Of course she made a decision that risked creating the tragic situation, but so did Vince when he decided to knock on her door (especially given that he has a murderous twin). But if Annie is neither at fault nor the main contributor to the problem, there is simply no way in which it is fair for her to bear the full burden of the situation. And so we cannot reasonably hold it against *her* that she attacked Vince in *Mistaken Attacker*.

The central motivation for this theory thus does not take fully seriously what I called the first-personal side of the question. The intuition that it

[12] See Quong, "Liability to Defensive Harm," 67–8. Quong introduces the evidence-relative qualification to deal with certain cases of risk-imposition. Since I do not see an important difference between the proposal below and that part of his account, I leave this issue aside. In what follows, I consider the theory's suitability for dealing with problems of uncertainty, which may not have been Quong's purpose.

[13] Quong, "Liability for Defensive Harm," 69ff.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

is no fairer for Annie to bear the burdens of the tragic circumstances of *Mistaken Attacker* than Vince bears this out. Once we recognize that Annie's evidence matters for what she ought to do, and objectivist reasoning is to be rejected, we cannot then continue to appeal to objectivist notions of fairness. We need to recognize how this point affects other parts of the story about liability as well. But when the objectivist notion of fairness is replaced by an evidence-sensitive notion, we have effectively replaced Quong's hybrid theory with a thoroughly evidence-sensitive one.

A second way of dealing with uncertainty is suggested by Helen Frowe. Like me, Frowe believes that "it is a mistake to build our accounts of permissible defence around knowledge that [people] cannot have."[14] Nevertheless, Frowe resists the thought that liability is sensitive to the evidence of defenders. Instead, she wants to "distinguish between the justness of inflicting a harm, and justness of the harm itself."[15] This leads to a hybrid theory in which (a) the permissibility of Annie's defensive force depends on the evidence available to her,[16] but (b) Vince's liability to this force depends on the facts as they objectively are.[17] Consider again *Mistaken Attacker*. For Frowe, it was morally permissible for Annie to harm Vince in self-defense because of her epistemic situation. But Vince was nevertheless not liable to her use of force because, as a matter of fact, he was not going to attack Annie.[18]

This view is deeply problematic. Suppose Frowe is right, and theories of permissible defensive force based on objective conditions are useless because they require people to have information they lack. What, then, could be the use of an objectivist theory of liability? Many of the epistemic problems that impair our ability to know whether defensive force is permissible also impair our ability to know whether others are liable to such force. This is true both *ex ante*, when we are trying to figure out the moral

[14] Frowe, "A Practical Account of Self-Defence," 250. She calls this the "myth of 'full and accurate' knowledge."

[15] Frowe, "A Practical Account of Self-Defence," 260.

[16] More precisely: [I]f 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" (Frowe, "A Practical Account of Self-Defence," 252).

[17] "[W]hether 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" (Frowe, "A Practical Account of Self-Defence," 252–53).

[18] Frowe's account aims to improve upon Kimberly Ferzan's earlier account in Ferzan, "Justifying Self-Defense." Ferzan was one of the first to draw attention to the problem of uncertainty in the context of liability to defensive harm. Ferzan correctly points out that a theory of liability that presupposes that defenders know things that they could never know (as most "objective" theories do) is useless or worse. However, Ferzan's own account insists that epistemic considerations only come into play once an objective condition of culpability on the part of the liable person is met. But of course the very same objection can be raised against this objective condition as Ferzan raises to objective theories more generally. For elaborations of this objection, see Frowe, "A Practical Account of Self-Defence"; and Jeff McMahan, "Self-Defense and Culpability," *Law and Philosophy* 24 (2005): 751–74.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

status of those with whom we are dealing, and *ex post*, when we are trying to figure out, say, whether to blame or punish people who have defended themselves. If such uncertainty is sufficient to render objectivist theories of permissibility useless, it is also sufficient to render objectivist theories of liability useless.

Frowe insists on a hybrid theory, instead of a fully evidence-based view, because she thinks that letting go of an objective standard of liability goes too far. After all, such a view entails that our moral protections against defensive harm would not just depend on facts about ourselves, such as what we have done, the position in which we find ourselves, and so on, but (also) on facts about others, such as their evidence, justified beliefs, and so on. That move, she thinks, is so uncomfortable as to be unacceptable.[19]

But a hybrid theory is no less uncomfortable. Even if our rights are not sensitive to others' evidence, the permissibility of others using defensive harm is. So the hybrid theory does not prohibit any uses of force allowed by a more thoroughly evidence-based theory. As a result, these theories render us equally vulnerable to others using force. And surely the source of discomfort is the possibility of attack, not the mere absence of rights against attack.

In fact, and more importantly, combining an objectivist standard of liability with an evidence-based standard of permissibility runs the risk of incoherence. The correlativity of claim-rights and duties means that if Vince has a right that Annie not use force against him, then Annie has a duty not to do so. To be nonliable to defensive force just is to have the right that others not expose one to such force. It follows that, other things equal, it is impermissible for Annie to use such force against Vince.

By separating these issues, Frowe's theory denies the basic correlativity of rights and duties. And we cannot deny that idea without radically altering our conception of rights in general, and accepting an idea of rights that lacks practical upshot or relevance. That is too high a price to pay for a theory of liability.

[19] In fact, Frowe does not even think that we should base permissibility on evidence in general. The context of self-defense is special, she writes, because self-defense "is by its nature urgent: it does not allow for the deliberation or investigation that we ought to require in other parts of morality" (Frowe, "A Practical Account of Self-Defence,"257). If someone points a gun at you, you have to make a snap decision about whether or not it is permissible to kill her in self-defense. This, too, is unconvincing. Decisions about defense are hardly unique in being time-sensitive. And decisions about defensive harm can allow for extensive deliberation. Real-world examples include preemptive and, especially, preventive wars. Given that there is no clear distinction between decisions about defensive force and other parts of life, the cost of postulating different sets of moral rules for these is considerable. For one, it would make it hard to see when rival (inconsistent) sets of moral rules apply, and thus what we ought to do. It would also make it hard to understand why one set of rules applies rather than another, or to what reasons our actions ought to be sensitive, and thus to solve uncertainty or disagreement about right action.

IV. Filling Out the Best Bet

The conclusion thus stands that the conditions under which people become liable include those in which treating them as liable is part of our morally best bet. This section begins the task of filling out this account of liability as part of one's best bet. What I will say does not exhaust the issue of liability or (*a fortiori*) the idea of a best bet, but identifies core cases in which people are liable.

The account I offer departs from two important Lockean commitments. In the beginning of the *Second Treatise*, Locke discusses the reasons why people can be justified in harming others, despite the fact that their rights normally protect them against this. Locke's answer took the form of a forfeiture argument. Sometimes, people can lose (forfeit) those rights and thereby become liable to otherwise prohibited harms. Locke's discussion offers two rationales for holding people to be liable in this way.[20] One is that people become liable when they threaten the rights of others. The other is that people become liable when they decide to wrong others.

The first of these is based on the moral imperative to preserve life — or, as Locke put it, "the Peace and *Preservation of all Mankind*."[21] Part of the point of morality is to ensure that people can live together safely. This aim lies at the foundation of our rights — which are to offer us key protections against being wronged or harmed by others — as well as their limits — including forfeiture. The rationale for having rights in the first place thus also allows for their enforcement. It offers the protection that our rights were meant to provide all along. We can forfeit rights, therefore, when our presence poses a threat to the rights of others.

The second rationale departs from the thought that one's decisions that would, if realized, foreseeably violate the rights of others can lower one's own moral standing, and thereby lower the rights-protections that our standing would otherwise warrant.[22] Let us call such decisions culpable decisions. We can become liable by making these culpable decisions because our status as rights-holders expresses, among other things, the fundamental moral equality of people. Rights offer protections against the unilateral imposition of ends by others. But this means that our moral status, and its accompanying protections, is importantly reciprocal. By demonstrating the lack of regard for others involved in culpably deciding to violate their rights, one undoes this reciprocity. Such culpable decisions

[20] Perhaps I should say that I find these rationales in Locke's discussion. I leave aside whether Locke would have endorsed these. My aim here is not Locke scholarship; it is an attempt at drawing insight from his text.

[21] John Locke, *Two Treatises of Government*, ed. P. Laslett (Cambridge: Cambridge University Press, 1988 [1689]), *Second Treatise*, sect. 7.

[22] Acting in a rights-violating manner entails acting without moral justification. Violating a right is different from infringing a right. The latter are normally in conflict with the demands of other people's rights, but are nevertheless done with moral justification.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

demonstrate that one does not consider oneself bound by the moral law, and thereby forfeit one's own rights-protection provided by that law.[23]

The two approaches to liability can come apart. It is possible for us to violate the rights of others without being culpable. If I bump into you because I am looking at my phone rather than where I am going, thereby knocking you off a cliff, I violate your rights. However, I would not be culpable in the manner described above, since I did not make a decision that foreseeably led to you falling down.[24] Nevertheless, I would be liable to your defensive force.

It is also possible for people to culpably act in ways that would violate another's rights without thereby succeeding. When the hit man takes aim and pulls the trigger, but his gun malfunctions, he is most definitely culpable for attempting to kill you. But it does not make sense to say that he violated your right to life. He made a decision that would, if realized, foreseeably violate your rights, but failed.[25] Nevertheless, in virtue of the second Lockean strand, the hit man is also liable to your defensive force.

Some might worry about the idea that people can become liable despite not actually posing a threat to others. However, it is not plausible that actual threats are necessary for liability.[26] For one, this would make the targets of culpable decisions highly vulnerable. The hit man adopted all the motivational parts requisite for performing a grave rights-violation. He put himself in a position where, as soon as the option of killing his victim becomes available, he will engage in precisely the sort of rights-violation around which the idea of liability above is organized. It would thus be unreasonable to demand of the hit man's target that she postpone initiating defensive force until the hit man gets his weapon back in order. At that point, it may well be too late.

Moreover, such a view has perverse implications. Suppose that the target, upon seeing what the hit man is trying to do, chooses to defend herself. If actual threats are required for liability, it would follow that once

[23] As Locke put it: "In transgressing the Law of Nature, the Offender declares himself to live by another Rule than that of *reason* and common Equity, which is that measure God has set to the actions of Men, for their mutual security; and so he becomes dangerous to Mankind, the tye, which is to secure them from injury and violence, being slighted and broken by him. Which being a trespass against the whole of Nature, every man upon this score, by the Right he hath to preserve Mankind in general, may restrain, or where it is necessary, destroy things noxious to them, and so may bring such evil on any one, who hath transgressed that Law, as may make him repent the doing of it, and thereby deter him, and by his Example others, from doing the like mischief. And in this case, and upon this ground, every Man hath a Right to punish the Offender, and be the Executioner of the Law of Nature" (*Second Treatise*, sect. 8).

[24] Thus, should you survive and be injured I ought to pay you compensation for the rights-violation. But it would not make sense for you to blame me for injuring you, as it was an accident. (It would make sense to blame me for my negligence.)

[25] He may, of course, still violate some other right of yours, such as the right not to expose you to the danger of being shot.

[26] As required, for example, by Quong's theory.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

the target takes out her own gun, aims at the hit man, and starts to pull the trigger, *the target* becomes liable. After all, she is posing a threat against the hit man who is not (on such an account) liable. This is clearly unacceptable. Suppose the hit man were to quickly repair his gun and shoot the target before she has a chance to shoot him. In that case, accounts that require actual threats for liability entail that the hit man would not violate the target's rights. After all, the target had become liable. This is bizarre. The hit man set out to kill his target in cold blood, yet it now turns out that killing her has become permissible.[27] Indeed, calculating killers might exploit this possibility. Thus, the hit man might himself disable the gun, knowing that he might be able to fix it quickly enough to outdraw and kill his target. This is a license to kill for unscrupulous wrongdoers who are willing to take the chance of being outdrawn by their victims.

We need, then, to bring together both the culpability and rights-violation approaches above. On the account I propose:

> *Y being liable to defensive force is part of X's morally best bet if X is justified in believing either that Y is about to perform an objectively rights-violating act, or Y has culpably acted in a way that would, if successful, be such a rights-violating act, or both.*

A few things are worth noting. First, and obviously, the proposal renders Y's liability sensitive to what X's best bet is in light of the evidence. Hence the role of X's justified beliefs.[28] However, because the account proposed here focuses only on what X is justified in believing about the dual sources of liability, we cannot simply move from this judgment to the broader judgment that it is X's best bet that the use of force is morally justified. What X is justified in believing about these conditions is only part of what goes into X's correctly calculated best bet about what to do.[29]

Second, the dual condition above talks about Y performing an *objectively* rights-violating act. Y performs such an act if Y's action would violate someone's rights, and thus be something that Y (other things equal) ought not to do, were Y to have all the relevant evidence available. Part of the

[27] One might object that after the hit man repairs his gun, he is again a threat to the target and thereby again becomes liable. But this is mistaken. The hit man repaired his gun *after* his target became liable, and thus does not pose a threat to a non-liable person.

[28] The idea of being justified in believing something is complex. For present purposes, I will consider what X is justified in believing and X's available evidence as more or less interchangeable. That is, I will consider X's being justified in believing something in terms of X having sufficient evidence concerning that thing. However, this is ambiguous between the evidence that is in principle available to X and the evidence of which X has in fact availed himself. I think the correct account is organized around the latter for reasons outlined by Zimmerman in chapter 3 of *Ignorance and Moral Obligation*. But the account of liability does not turn on this, as it is a general concern about how to treat the evidence-relative dimension of morality. For this reason, I sideline these complicated issues.

[29] I do not claim, for instance, that X's best bet is also made up of X being justified in believing that force is necessary, proportionate, and so on.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

proposed dual account thus ties Y's liability to X's evidence about whether Y is to perform such an act.

This objectivist-sounding element prevents certain problems that would plague an even more thoroughly evidence-sensitive account. Consider a variation of an objection offered by Quong to such thoroughly evidence-sensitive accounts. Suppose a group of "duped soldiers" have been given convincing but false evidence that some neighboring group is about to attack them. If the soldiers set out to attack their neighbors in (mistaken) self-defense, and if the soldiers could not be liable unless they do something that (other things equal) they ought not to do in light of their evidence, then the duped soldiers will be non-liable. This is problematic since the neighbors are clearly permitted (other things equal) to defend themselves against the duped soldiers.[30]

Quong is right that we should stay away from such a thoroughly evidence-sensitive view. By focusing solely on the epistemic circumstances of agents, such a view loses sight of what ultimately matters: avoiding objectively rights violation acts. This is the reason the condition above refers to such objective rights-violations. On that condition, the duped soldiers do become liable. By attacking their neighbors, they are about to do something that would violate the neighbors' rights, were they to have all the evidence, and the neighbors know this.[31]

Third, the proposal entails that for Y to be liable to X, it must be true that X has evidence that the target of Y's action is itself not liable to Y's defensive force. This need not always be the case. If X is attacking Y, and X has evidence that Y has a right that X not attack Y, and X has evidence that Y is about to engage in permissible self-defense, then X will not be justified in believing that Y is about to violate X's objective rights. For X's evidence is that X is liable to Y's use of defensive force, and therefore that Y is not about to violate X's objective rights.

Finally, consider how this account of liability applies to cases like *Mistaken Attacker*. There, Annie (mistakenly) attacked Vince because of false but convincing evidence that Vince was about to commit an objective rights-violation. It follows, then, that Vince was liable to Annie. However, it *also* follows that Annie is liable to Vince's use of force. For, despite her evidence, Vince does not satisfy the Lockean conditions, and Vince knows this. Thus, when Vince notices that Annie is about to attack him, his evidence suggests that Annie is about to engage in an objective rights-violation. As a result, Annie is liable to Vince's self-defense.

[30] Quong discusses a slightly different case because he considers a slightly different evidence-based view from the one I have proposed here. Nevertheless, his discussion contains the basic building blocks of this (for my account more threatening) objection. See Quong, "Liability to Defensive Harm," 62.

[31] The same is true, of course, for relevantly informed third parties. With respect to them, too, the duped soldiers are liable to defensive force. Note that, in these respects, the case of the duped soldiers is equivalent to *Mistaken Attacker*.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

Upon reflection, this strikes me as the correct verdict. Since it is true of neither Annie nor Vince that they ought not to attack the other, and since neither Annie nor Vince is in the relevant sense primarily responsible for the situation in which they find themselves, both Annie and Vince ought to be liable. *Mistaken Attacker* constitutes a tragedy in which each, through no fault of his or her own, becomes liable to the other's use of force.

V. Objections

In closing, I will discuss some objections to the account outlined above. I will first discuss two objections to the view's evidence-sensitivity, then an objection to how the evidence-sensitive and rights-based parts fit together.

The most obvious worry about the account's evidence-sensitivity is that it implies that if X has faulty evidence about Y, then Y can become liable to X even though Y is in fact innocent. I admit that this is both counterintuitive and worrying. Indeed, I used to think it was clearly wrong myself. However, there are at least two reasons why, upon closer inspection, what may seem counterintuitive or worrying is not so hard to swallow. The first and most important reason is the one we have already seen: the price of rejecting these implications is even higher. For the denial of this basic point requires either the rejection of the correlativity of rights and duties, or an unacceptable implication in *Uncertain Rescue*.

The second reason is that things are slightly better than they might seem. For even though the proposed theory allows that Y can become liable to X's defensive force because of X's faulty evidence, it also implies that if Z has better evidence, then Y is not only non-liable to Z but X is liable to Z as well. This means Z is permitted to intervene on Y's behalf — as, of course, is Y himself. Consequently, if people are generally relatively likely to evaluate correctly the liability of others, the theory will in practice strongly tend toward the actual enforcement and preservation of rights.

Another objection holds that the theory I have proposed does not allow enough room for cases in which people act wrongly — for example, by violating the rights of others — but ought to be excused. The account above seems to imply that for people to violate the rights of others, they must act contrary to their evidence. And doing that seems simply inexcusable. But surely, the objection goes, it is possible to do the wrong thing even though we should be excused. Thus, an evidence-based account is to be rejected.[32]

This is a mistake. There can be cases in which people do not do what is, in light of the evidence, their best bet, and thus (other things equal) violate the rights of others, but should be excused. For it can be reasonable, or at least understandable, to believe things that we are not justified in believing. One example is when we act on the basis of beliefs that result from certain

[32] For a variation of this objection, see Ferzan, "Justifying Self-Defense," 713 and 729.

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press

standard cognitive biases. When we suffer from biases in ways that we are not culpable for failing to correct, we can believe that we are acting permissibly in ways that ought to be excused. Since the account above refers to our justified beliefs, it does not preclude acting wrongly in ways that ought to be excused.

A third and final worry concerns the compatibility of the proposed theory's rights-based and evidence-sensitive parts. Most worryingly, perhaps, the account allows people to gain or retain rights that, in objectivist terms, they should lose. Indeed, it allows for this to happen *as a result* of their wrongdoing. All that is required is for the evidence to line up in the requisite way. This opens up worrying possibilities of manipulation. Suppose that an evil group is in the process of killing many innocent people. And suppose that they try to hide the evidence of their heinous deeds, say by threatening to kill the families of anyone who would report their killings. If the evil group were successful, it would follow that the group is not liable to defensive force.

For a theory organized around the moral importance of rights, this looks like a problem. Indeed, it is especially a problem for the theory defended here, since it bases liability in part around Lockean motivations. Rights-holders, I said, enjoy a kind of moral status or standing, and so this theory of liability implies that people can acquire moral standing through wrongdoing. The evil group, after all, manages to retain its right against defensive force through unjustly threatening others. This is the exact opposite of the moral reciprocity discussed above.

Obviously, this is a worrying result. But in the end I do not think it poses a problem for the theory of liability defended here. For it is in general a possibility that people acquire rights despite and through their wrongdoing. Suppose that Beth's will states that, were she to die, Chuck inherits all her money. And suppose Chuck murders Beth. Chuck would thereby gain a right to Beth's money through his wrongdoing. Chuck would thus acquire the standing of a right-holder through his own wrongdoing. This worrying possibility must be allowed in general, it seems. As a result, its existence does not constitute an objection unique to the theory I have outlined. Unfortunately, morality sometimes allows the especially heinous to manipulate their standing.

Philosophy, University of North Carolina, Greensboro

https://doi.org/10.1017/S0265052516000133 Published online by Cambridge University Press