

The Toss-Up Between a Profiting, Innocent Threat and His Victim*

Susanne Burri
s.burri@lse.ac.uk

3 July 2014

1 Introduction

Imagine that, through no fault of your own, you find yourself at the bottom of a deep well.¹ Thugs have picked up an innocent person—call him Bob—and have thrown him down the well. Bob is now falling towards you. If you do nothing, your body will cushion Bob’s otherwise lethal fall. This will guarantee his survival, but it will kill you. If you shoot your ray gun, you vaporize and kill Bob, thereby saving your life. Are you morally permitted to shoot your ray gun?

If you are, then it is permissible to kill a *profiting, non-responsible threat*. Bob is a *threat* (as opposed to an attacker) because he endangers your life through no immediate action of his own. He is *non-responsible* because he neither intended to become a threat, nor was becoming a threat a consequence of an intentional action of Bob’s that he could reasonably have been expected to foresee.² Bob *profits* from you because he depends

*This paper is forthcoming as an article in the *Journal of Political Philosophy*. For valuable feedback and discussions I am grateful to Brian Berkey, Alexis Elder, Jeff McMahan, Michael Otsuka, Jonathan Parry, Jonathan Quong, and an editor as well as two anonymous referees of the *Journal of Political Philosophy*. Earlier versions of this paper were presented at the Warwick Graduate Conference, the Yale/UCConn Graduate Conference, the London-Berkeley Graduate Conference, the Brave New World Postgraduate Conference, and at a Choice Group Seminar at the London School of Economics and Political Science.

¹This hypothetical situation is modeled on an example in Nozick, 1974, p. 34.

²This follows the use of the term ‘moral responsibility’ in McMahan, 2009, pp. 167-8.

on your presence for his survival.³ There is no agreement among moral philosophers whether self-defence against a profiting non-responsible threat is permissible. In this paper, I argue that the right thing to do when faced with such a threat is to flip a coin. By flipping a coin to determine who should live, you ensure that the indivisible good of continuing one's life is distributed as fairly as is possible without wasting it.

The structure of this paper is as follows. In sections 2 and 3, I present the main arguments that have been brought forward in favour of and against the moral permissibility of killing a non-responsible threat in self-defence. In section 4, I argue that all of these arguments rest on a misconception, and that flipping a coin emerges as the right thing to do once this misconception is corrected. In section 5, I compare and contrast a number of cases to test the plausibility of the arguments brought forward earlier. In section 6, I argue against the idea that a potential victim has a prerogative to save her own life.

2 Arguments in Favour of Permissibility

Judith Thomson (1991, esp. pp. 300-3) argues that killing a non-responsible threat in self-defence is permissible because a non-responsible threat loses his right not to be killed once he starts threatening the life of his victim. Thomson's idea is that Bob loses his rights against you when he starts falling towards you in a life-threatening manner *without your having threatened his life first*. Given that Bob started falling towards you in this unprompted way, Thomson believes that Bob would *violate* your right not to be killed if he killed you. To Thomson, you have a right to enforce your right not to be killed. This is why she concludes that Bob must have lost his right not to be killed when he started falling towards you.

According to Thomson's analysis, moral agency is not a necessary condition either for forfeiting one's rights or for violating someone else's. Bob's

³In this paper, my primary concern is with *profiting* non-responsible threats. Unless stated otherwise, all references to non-responsible threats should thus be read as references to *profiting* non-responsible threats. In the literature, the distinction between profiting and non-profiting threats is not usually drawn. One exception is Victor Tadros, who mentions the distinction, but thinks it irrelevant for his arguments. See Tadros, 2011, p. 248. For the remaining authors I mention in this paper, I will simply assume that they mean their arguments to apply equally to both profiting and non-profiting non-responsible threats.

falling towards you is not something he is in control of or responsible for; it is something that merely—and through no fault of his own—happens to him. Thomson argues that it nevertheless grounds Bob’s liability⁴ to defensive harm by making it the case that Bob is threatening to crush you, which for her means that he is threatening to violate your right not to be killed.

Kamm (1992, pp. 45-55; 2012, p. 61)⁵ argues that Bob would have to suffer greater harm to avoid harming you than an innocent bystander would have to suffer to keep Bob from harming you. To give an example, Kamm (1996, p. 96) thinks that if Bob could redirect himself away from you and save your life at the cost of breaking his leg, then he would have to redirect himself, even though an innocent bystander would not have to break one of his legs to keep Bob from harming you. To Kamm (*ibid.*), if Bob is unable to impose on himself the harm that he would be duty-bound to suffer for your sake, others may impose it on him on his behalf. But why does Kamm think that Bob’s moral standing *vis-à-vis* you is weaker than that of an innocent bystander? Kamm (1992, p. 47) claims that Bob’s falling towards you means that he is “is headed to where [he] should not be, on the victim, who has a right not to be occupied.” To Kamm (*ibid.*), “[one] simply has a right not to have someone on the body or property to which one is entitled, even if the wind put them there.” In Kamm’s view, Bob is thus liable to defensive harm not because he threatens to violate your right not to be killed, but because he threatens to violate your right not to be occupied. If either the victim or the threat will be harmed very seriously, Kamm (1992, p. 54; 2012, p. 61) claims that it is permissible to side with the victim, presumably because the threat, but not the victim, is in a morally inappropriate position.

Like Kamm, Victor Tadros (2011, pp. 248-56) believes that Bob does not threaten to violate your right not to be killed, but that he would nevertheless have to suffer greater harm for your sake than an innocent bystander would have to suffer. To Tadros, this is so because we have a special responsibility for what our bodies do, even when what our bodies do is not a product of our agency. Tadros furthermore argues that you may impose more harm on a non-responsible threat than he has an enforceable duty to suffer. To Tadros (2011, p. 148), the harm someone has an enforceable duty to suffer is the harm that we may intentionally inflict on them when manipulatively using

⁴The language of liability was popularized by McMahan. Thomson does not use it, though she is in effect relying on the same notion.

⁵See also Kamm, 1987, and Kamm, 1996, esp. p. 96.

them as a means to promote some further end. But if you vaporize Bob, you do not use him as a means; you merely eliminate him. Tadros (2011, p. 244) thinks that eliminative harm is easier to justify than manipulative harm, as only the latter involves seeing people as “material for use to promote the greater good.”⁶ Given his idea that Bob would have to lose a foot to keep himself from crushing you, Tadros (2011, p. 251) then argues that once we “accommodate [the] difference between manipulative and eliminative harming by [increasing] the cost that would have to be borne by the threatening party”, it becomes plausible that you may kill Bob eliminatively.

3 Arguments Against Permissibility

In contrast, Jeff McMahan and Michael Otsuka argue that a non-responsible threat is not liable to defensive harm because he hasn’t done anything that would justify a diminution in his moral status.⁷ Because they think it impermissible to kill someone who is not liable to defensive harm in order to save a single person, they conclude that it is impermissible to kill a non-responsible threat in self-defence.⁸

McMahan (1994; 2002; 2005) argues that we become liable to lethal defensive harm whenever we responsibly create a “forced choice between lives” (McMahan, 2002, p. 403), or bring it about that either we or someone else must die. Succinctly put, McMahan conceives of “the morality of self-defence as a matter of justice in the distribution of harm” (McMahan, 2002, p. 402), and maintains that it is just to defensively impose harm on someone only if that person is morally responsible for making its distribution unavoidable. It follows that according to McMahan, it is unjust—and hence impermissible—to kill a non-responsible threat in self-defence, as such a threat by definition lacks moral responsibility for his threatening status, and is thus not responsible for making the distribution of harm unavoidable.

⁶Tadros notes that his distinction between manipulative and eliminative agency coincides with Warren Quinn’s distinction between opportunistic and eliminative direct agency. See Quinn, 1989, p. 344.

⁷By the ‘moral status’ of a human person I mean the moral standing she enjoys vis-à-vis others. Someone’s moral status is higher the less we may do to her against her consent, and the more we must do for her at some cost to ourselves or others.

⁸Kai Draper (1993) was an early defender of a view similar to McMahan’s and Otsuka’s.

Otsuka (1994) starts from the observation that it is impermissible to “ride roughshod”⁹ over an innocent bystander to save one’s own life. To illustrate, imagine that a piano is falling towards you, and that the only way to prevent it from crushing you is to vaporize it with your ray gun. The problem is that if you shoot the piano, you will also end up vaporizing Paul, an innocent stranger who is standing at one of the windows of the building from which the piano is falling. May you vaporize the piano and ride roughshod over Paul to save your life? Otsuka thinks it intuitively highly plausible that you may not. He argues that if this is right, it should also be impermissible to vaporize Bob, as there is no moral difference between Paul and Bob. According to Otsuka, this can be seen by constructing an intermediate case in which Paul no longer stands at the window, but sits—through no action that would make him responsible for falling towards you—on top of the falling piano. If you vaporize the piano, Paul dies; if you refrain from using your ray gun, Paul survives unscathed, and you die. Otsuka would maintain that shifting Paul’s location from ‘standing at the window’ to ‘sitting on top of the piano’ cannot make Paul liable to defensive harm. Put differently, Otsuka denies that it can make a moral difference whether an innocent person is located next to a threatening object or somehow ‘tied up’ with it. Moving from the intermediate case to the non-responsible threat case, Otsuka removes the piano from under Paul, so that it is now Paul’s own weight that threatens to crush you. Otsuka would again assert that no moral difference is involved in this move. To him, the fact that it is now Paul’s own body that threatens to crush you cannot render it permissible to vaporize Paul when before it wasn’t.

⁹This is Thomson’s terminology. See Thomson, 1991, p. 290.

4 The Toss-Up Between a Profiting Innocent Threat and His Victim

I believe that the arguments presented in the last two sections are misguided. Here I explain why, and introduce what I take to be the right way of looking at things.

4.1 The Misconception

The positions presented so far all identify a purportedly *decisive moral asymmetry* between Bob and you. Given this, it seems likely that they are based on considerations along the following lines: “It is either permissible to kill a non-responsible threat in self-defence, or it is impermissible. Either way, there has to be a moral asymmetry between the threatened person and the threat, an asymmetry that favours either the victim, in which case the killing is permissible, or the threat, in which case the killing is impermissible. It is our task as moral philosophers to identify and explicate this asymmetry.”

Thomson, Tadros, and Kamm identify an asymmetry that favours you. Though labeled ‘non-responsible’ for threatening you, Bob is nevertheless characterized as morally compromised in some way: he threatens to violate your rights, finds himself in the wrong position vis-à-vis you, or is responsible for the harm his *body* threatens you with. Bob’s inappropriate relation to you then grounds his liability to defensive harm, and makes it permissible for you—the *victim*—to defend yourself.

McMahan and Otsuka identify an asymmetry that favours Bob. His falling towards you merely happens to him; it is not something that he is in control of. If he ends up harming you, he cannot reasonably be held responsible for doing so. The same cannot be said of you. If you shoot your ray gun, you do so voluntarily and deliberately, as a fully responsible moral agent. Given his innocence, Bob cannot be liable to defensive harm, hence you—the *morally responsible agent*—mustn’t shoot your ray gun.

Neither way of looking at things is wholly implausible. However, in explicating various asymmetries, both miss the fundamental *symmetry* that in my view characterizes the situation. Notice how similar Bob and you really are. Bob is innocent; he cannot be blamed for falling towards you. But you too are not to blame for the predicament you find yourself in! You

did not know, and could not possibly have known, what you were getting yourself into when you stepped into the well. Yet as things are, you can save yourself only if you vaporize Bob, and Bob will survive only if he ends up crushing you. Both Bob and you faultlessly find yourselves in perilous situations from which you will escape only if an innocent other dies through you. Given this fundamental symmetry, it seems promising to frame the problem of the permissibility of killing a profiting non-responsible threat in self-defence in *distributive terms*, and to think of the threat and his victim as individuals with competing claims to the important good of continuing one's life (cf. Levine, 1984).

Insofar as Bob and you are similar, your claims to the good of a continued life tie. But it would be rash to conclude that your claims are all things considered equal, as there are important differences between the two of you also. In the next section, I examine these differences, and discuss whether they strengthen or weaken the non-responsible threat's position vis-à-vis his potential victim.

4.2 Accounting for the Differences Between a Profiting Innocent Threat and His Victim

A first important difference between Bob and you is that you act towards Bob as a morally responsible agent, whereas whatever Bob does to you he does to you merely qua falling body. McMahan and Otsuka believe this asymmetry weighs crucially in Bob's favour. Their idea is that Bob is not liable to defensive harm because he hasn't done anything that would justify a reduction in his moral status. Given that Bob's threatening you is not something he is responsible for or in control of, Bob cannot sensibly be said to be under a duty not to threaten you, and you cannot accordingly have a right against him that he not threaten you. Bob, on the other hand, has a right against you that you not kill him, and you have a duty not to kill him. From Bob's perspective, this way of looking at things appears plausible: surely what Bob *ought* to avoid, he has to be *able* to avoid?

Thomson disagrees. Her idea is that we can threaten to violate a right, and thus become liable to defensive harm, merely by being turned into a threat. From Bob's perspective, this way of looking at things is less compelling than McMahan's and Otsuka's. Why should Bob accept a reduction in his moral status because of something he is neither responsible for nor

able to control? We can make better sense of Thomson's view once we give up the perspective of the threat, and consider the problem from the victim's point of view. From your perspective, your rights to non-interference enable you to lead a self-determined life in accordance with your own interests by giving you "primary say over what may be done to [you]" (Quinn, 1989, p. 309). Rights to non-interference make you *sovereign* over your own body: they make it the case that other people are *trespassing* on your territory, and may therefore be fought off, whenever they cross the protective boundary that your rights to non-interference draw around you. As far as his impact on you is concerned, Bob's lack of responsibility is irrelevant. His body threatens to damage yours, and unless you are permitted to fight Bob off, this will have a negative effect on your ability to lead your life as you please (Thomson, 1990, pp. 222-3). If we think of your rights to non-interference as drawing a protective boundary around you that you have a right to enforce, Bob is *unlucky* to be crossing that boundary entirely involuntarily and non-responsibly, but that doesn't change the fact that he is crossing it.

I believe that Thomson's view contains an important truth. From your perspective, if your rights to non-interference become ineffective against Bob while his remain effective against you, Bob's bad luck of being turned into a threat is passed on to you, and this seems unjust. But Thomson's view has the contentious implication that 'ought' need not imply 'can', or that Bob continues to be under a duty not to violate your rights even when he has temporarily lost control over his actions. This strikes me—and many others—as unacceptable. Kamm's view faces a similar problem. Like Thomson, Kamm focuses on the victim's rights to non-interference, and insists that these rights do not become wholly ineffective against a non-responsible threat. If we are committed to the idea that 'ought' implies 'can', do we thus have to agree with McMahan and Otsuka that Bob is not liable to defensive harm?

Tadros does not think so. He suggests we instead drop the assumption that a person becomes liable to defensive harm only if he threatens to *violate* another's rights. To Tadros, Bob is liable to defensive harm because we are all "responsible for what [our bodies do], even when that is not a product of [our] agency". simply because our bodies are—at least to some extent—what *we* are (Tadros, 2011, p. 255).

What should we make of this idea? For McMahan and Otsuka, it might

simply beg the most important question. McMahan and Otsuka stress the fact that Bob is a *non-responsible* threat; Tadros in essence denies that a human person could be such a thing. McMahan and Otsuka might well agree that we are responsible for what our bodies do to the extent that we are able to control them, or are morally responsible for not being able to control them. But why should—or how could—our responsibility for our bodies extend beyond that? Suppose I own a piano which, for reasons beyond my control, crushes you. How am I morally responsible for the harm the piano causes *merely* because the piano is mine? Given this, I don't think we should agree with Tadros that Bob is responsible for what his body does. But then what are we left with? Is the fact that Bob poses a threat to you as morally irrelevant as McMahan and Otsuka claim it is?

Kamm and Tadros both base their arguments on the following observation: *if* Bob *could* avoid harming you, he would have to suffer substantial harm to do so. This seems to me importantly right. If Bob could redirect himself away from you to avoid crushing you, he would have to do this even if it would cost him a foot. But why would Bob have to sacrifice his foot? In essence, because our morality is built around stringent rights to non-interference, or on a strong presumption against harming innocent others. As long as we are able to avoid harming innocent others, we generally have to go to great lengths to do so. The mere fact that we lack responsibility for finding ourselves in a situation where we pose a threat to innocent others does not make it permissible to harm them. This is why Bob would have to suffer substantial harm to avoid harming you if he could, or why Bob would be violating your rights to non-interference if he crushed you when he could have redirected himself away from you at reasonable cost to himself.

Moral rules that are built on a strong presumption against harming innocent others neither allow us to minimize the harm that comes to innocent persons, nor do they allow us to distribute harm equally among similarly undeserving recipients. But they shield us from harm that is imposed on us by other people as long as we do not inflict harm on others. In this way, they provide us with the peace of mind that other people must not, in general, threaten our physical integrity. Knowing that our bodies are largely off-limits for those around us allows us to live together in peace while pursuing our own interests.

Yet even if Bob *would* have to make great efforts to avoid harming you if

he *could*, this seems to be no more than a hypothetical consideration. After all, we have assumed that Bob cannot actually avoid harming you. You, on the other hand, can avoid harming Bob. Given this, and given our morality's strong presumption against harming innocent others, doesn't it follow that given the way things actually are, you ought to refrain from shooting Bob?

To Tadros and Kamm, the reason why Bob would have to avoid harming you if he could applies also when he cannot in fact avoid harming you. If Bob were responsible for whatever his body did, or if you had a right against Bob not to have Bob's body on yours, this would hold true irrespective of Bob's ability to control his fall. But I have argued that Tadros and Kamm's ideas are not compelling. Contrary to what Tadros and Kamm suggest, moral rights to non-interference are best thought of as held symmetrically among moral agents. Qua moral agent, you have rights to non-interference against other moral agents, and they have the same rights against you. Only if one moral agent violates the rights of another is this symmetry disturbed so that balance can be restored. The protective logic of our moral rights to non-interference is thus built around—and for its functioning depends on—the idea of mutual recognition between moral equals.

But note that if this is right, Bob retains his rights to non-interference when he starts falling towards you not only because he lacks responsibility for falling towards you. Lack of responsibility for threatening another is a necessary condition for keeping one's rights to non-interference, but it is not sufficient. Bob also retains his rights because he retains his moral capacities, or because he remains the kind of creature who would have to avoid harming you if he could. In other words, if you must not vaporize Bob's out-of-control body, this is partly because Bob is the kind of creature who can understand that if he could avoid crushing you, he would have a duty to do so. But if you must constrain yourself towards Bob and let his body harm you *partly because Bob would (right now!) have to suffer substantial harm to avoid crushing you if he could*, it would be unjust if you could not impose that harm on Bob on his behalf. By analogy, suppose that as a parent, you ask your son to share his Halloween candy with his sister. You tell him that sharing is appropriate partly because his sister would also have to share her candy if she could. If so, it would be unjust and illogical if you made your son share his candy, but forbade him to help himself to his share of his sister's candies that happened to be out of his sister's reach but well within

his.

But what about the fact that whatever harm you impose on Bob, you impose on him as a morally responsible agent? Don't you violate Bob's rights if you harm him? As long as you inflict no more harm on Bob than he is duty-bound to suffer, you do not violate his rights to non-interference. Whatever harm Bob is duty-bound to suffer, he would—as a morally responsible agent—have to inflict on himself if he were able to. There is thus a certain amount of harm that may rightfully be assigned to Bob, given the situation he is in. Just as you may shove away a bystander who is unwilling or unable to move out of the way when you are running for your life, you can impose on Bob the more substantial harm he is duty-bound to suffer without violating his rights (cf. Overland, 2011, p. 550). Having said that, burdening Bob with more harm than he is duty-bound to suffer does not seem justifiable, and would therefore violate his rights to non-interference.¹⁰

So far, I have argued that there is a justice-based reason to make a non-responsible threat share in the costs that his out-of-control body threatens to impose on his victim. To the extent that a non-responsible threat would have to suffer harm to avoid harming his victim if he could do so, we may redistribute harm from the victim to the threat. This does not imply that Bob is liable to defensive harm. It merely means that his rights to non-interference are permissibly infringeable to the extent that there is a justice-based reason to make him bear part of the cost that he would otherwise impose on his victim.

If what I have argued so far is correct, where does it leave us with respect to the first asymmetry between Bob and you, i.e. the fact that whatever harmful effects you have on Bob are due to your responsible agency, whereas whatever harmful effects Bob has on you are not attributable to his responsible agency? Compared to a responsible threat, the asymmetry

¹⁰Tadros claims that you are permitted to kill Bob because you would not be using Bob as a means. I am not convinced by Tadros' argument in favour of this claim. His argument relies on the assumption that there are significant differences between two types of intentional agency, so that eliminatively killing someone is no worse than manipulatively cutting off their foot. At the same time, it is a consequence of his argument that liability to defensive harm does not vary with the type of responsibility that grounds it. According to Tadros, vicious threats, responsible threats, and non-responsible threats are all similarly liable to defensive harm. But if 'mode of agency' is a significant consideration in one context, why is it not significant when it comes to a person's liability to defensive harm?

weighs in Bob's favour, but less decisively than McMahan and Otsuka have argued. Though you would be permitted to kill a responsible threat in self-defence, you are not similarly permitted to vaporize Bob to save yourself. Instead, you have to take Bob's lack of responsibility into account and let Bob's right to non-interference constrain you. But Bob's moral status is not thereby equal to that of an innocent bystander. If Bob could avoid crushing you, he would have to suffer substantial harm to do so. This gives you a justice-based reason to make Bob share in the harm that he is threatening you with.

But Bob differs from an innocent bystander not only because he is threatening you. Recall the innocent bystander Paul, who is standing at one of the windows of a building from which a piano is falling towards you. Imagine what would happen to Paul if we removed you from the scene. Without you present, Paul could stand at the window and watch the piano fall without any risk of harm. Bob's situation is different. If you weren't there to cushion his fall, Bob would die on impact. Paul is thus merely a bystander to a dangerous situation not only because he does not threaten you, but also because his life is not at risk unless you involve him in your problems. Bob, on the other hand, depends on you for his life. Now by assumption, Bob's depending on you for his life is not something he is morally responsible for. But this does not suffice to make it morally irrelevant. After all, Bob's threatening you is also not something he is morally responsible for, yet it still serves to morally distinguish him from an innocent bystander. But why—and how—should Bob's profiting from your presence affect his moral standing vis-à-vis you?

Thomson (1971, pp. 48-55) has famously claimed that if a world-class violinist's life were to depend on his being hooked up to your circulatory system for the next nine months, you would be permitted to unhook yourself, thus killing him. To Thomson, you have no duty to sacrifice nine months of your life to avoid killing the violinist. Instead, she argues that his non-responsible dependence on you weakens his moral standing vis-à-vis you, so that you may do more to him than you would be permitted to do to an innocent bystander.¹¹ Kamm (1992, pp. 20-45) argues that the violinist's

¹¹This is not to say that responsibility is irrelevant to the assessment of dependence relations. If Bob were responsible for becoming dependent on you, his moral standing vis-à-vis you might well be more unfavourable than it is under the current circumstances, where he depends on you entirely non-responsibly.

dependence on you weakens his moral standing vis-à-vis you because he cannot be worse off through you than he would have been in your absence. If you unplug the violinist, he dies; but in your absence, he would have died anyway. Because of this, unplugging the violinist resembles letting him die. While you may have to shoulder significant costs to avoid letting an innocent other die, these costs are not usually thought to be as substantial as the costs that you have to shoulder to avoid killing an innocent other (Kamm, 1992, p. 71). Applied to the situation between you and Bob, if you vaporize Bob, you kill someone who would have died in your absence. By contrast, if you vaporize Paul, you kill someone who would have lived in your absence. This distinguishes Bob from an innocent bystander such as Paul.

The fact that a profiting threat depends on his victim not only distinguishes him from an innocent bystander; it also distinguishes him from his victim. Bob is dangerous—a threat—to you, whereas you usefully cushion Bob’s fall at best, or do not substantially alter his predicament at worst (if you vaporize him, he dies, but without you there, he would have died anyway). From a self-interested perspective, Bob thus welcomes your presence, whereas you cannot be said to welcome his. Following Thomson and Kamm, I think that this difference strengthens your moral standing vis-à-vis Bob. It does not give you license to treat Bob in any way you like, but it does seem to give you an advantage over Bob. Given that you are an innocent other who improves Bob’s chances of survival, whereas Bob is an innocent other who worsens your chances of survival, Bob should have to make an extra effort to avoid harming you if he could.

4.3 Weighing Claims

So far, I have argued that there is a justice-based reason to impose on Bob the substantial harm that he would be duty-bound to suffer for your sake. Does this shed any light on the relative strengths of Bob’s and your claim to the good of a continued life?

Suppose first that Bob was a non-profiting threat. I have argued that given our morality’s strong presumption against harming innocent others, Bob would have to suffer *substantial* harm to avoid harming you if he could. But would he be required to subjugate his own interests entirely to yours? I don’t think so. But why not? Given a strong presumption against harming

innocent others, why shouldn't Bob have to do whatever it takes to avoid killing you? Recall the innocent bystander Paul, whose life you are required to spare at the cost of your own. It may be very demanding to let yourself be crushed by a falling piano when you could in principle vaporize it. Yet given our morality's strong presumption against harming innocent others, you are at most *excused* if you vaporize the piano and kill Paul. But if you have to avoid killing Paul at the cost of your life, why wouldn't Bob have to avoid killing you at the cost of his life?

The important difference is this. You avoid harming Paul as long as you do not harm anyone, whereas Bob avoids harming you only if he inflicts harm on himself. Because Bob has to inflict harm on himself to avoid harming you, Bob can at most be required to share equally in the harm that his body threatens to impose on you. If he were required to pick up the lion's share, this would deny him the right, qua moral agent, to regard harm that he imposed on himself as on a par with harm that he let come to you. Given Bob's innocence, this would be an alienating requirement. Even if our morality is built on a strong presumption against harming innocent others, when avoiding harm to an innocent other means inflicting that harm on our innocent selves, we must not be required to treat ourselves as below par. Such a requirement would be inconsistent with a proper sense of self-worth.

Should Bob be permitted to give *more* weight to his own interests? If Bob harms you, he merely *allows* himself to harm you, and he merely *foresees* that harm will come to you. Allowing harm is often thought less bad than actively inflicting it, and merely foreseeing harm is often thought less bad than intending it. Does this mean that Bob should not have to intentionally inflict as much harm on himself than he lets come to you? Suppose you are the front passenger in a car whose driver starts fainting just as the car comes up behind a pedestrian. If you do nothing, the car will hit the pedestrian. Is the harm that you are required to suffer to avoid hurting the pedestrian reduced by the fact that whatever harm that you let come to him is neither intended nor actively inflicted? If you had already taken the wheel and started steering when you noticed the pedestrian, would this increase the harm that you would be required to suffer to avoid hurting the pedestrian? My hunch is: *not by much*. Why could that be so? It seems to me that our duties to avoid harming others are requirements to regard their bodies as territory that we need to steer clear of in the pursuit of our own interests.

The main failure of respect involved in harming an innocent person is *not regarding that person's body as off-limits for one's purposes*. This holds true for all modes of harming, and it makes them all similarly bad. In other words, even if Bob's mode of harming you is less disrespectful than other modes, Bob's attitude towards you is nevertheless very inappropriate if he fails to redirect himself away from you at moderate cost to himself.

To what extent might Bob's relatively respectful mode of harming you work in his favour? Out of two comparable options, Bob should probably be allowed to pick that option which slightly favoured his interests over yours. If one option gave both of you an equal chance of survival, and a second option improved Bob's chances (while reducing yours) by a couple of percentage points, Bob should probably be permitted to go with the second option. But note that we have not yet factored in the consideration that Bob is a *profiting* non-responsible threat. Taken by itself, the fact that Bob benefits from your presence cannot be of overriding importance. Whether other people's mere presence happens to be beneficial or burdensome to us is mostly a matter of our good or bad luck. Having said that, whether a person's presence is beneficial or burdensome is not entirely morally insignificant, as it can change the extent to which harming that person is disrespectful. In this way, it seems to me that a profiting threat would act unduly disrespectfully if he did not treat his victim's interests as on a par with his own. Suppose that as a profiting threat, Bob could again choose between two options, the first of which gave both of you an equal chance of survival, and the second of which was slightly skewed in his favour. Further suppose that Bob argued as follows: "I must not be required to pick the first option. That option is maximally demanding. It is the upper limit of what can be required of me without undermining my self-worth. Given that I do not trespass on my victim's territory in an especially disrespectful way, I should be cut some slack." It is not a profiting threat's place to argue in this manner. For even if Bob has to pick the maximally demanding option, he still has a 50% chance of survival, whereas if it weren't for you, he would simply die. Also, even if Bob has to pick the maximally demanding option, there is still a 50% chance that you will die, whereas in Bob's absence, your life would not be at risk. Given that circumstances have put Bob in such a favourable situation vis-à-vis you, Bob's attitude towards you should be one of humility. At any rate, he should not appeal to otherwise legitimate but relatively weak

considerations to discount your interests relative to his own. It follows that a profiting non-responsible threat is morally required to treat his victim's interests as on a par with his own. From a distributive perspective, this allows us to conclude that Bob and you have similarly strong claims to the good of a continued life.

4.4 The Distributive Perspective

When different people have similarly strong claims to a good, it is fair to assign equal amounts of it to each of them (Broome, 1984; Broome, 1991). In non-responsible threat cases, however, the good in question is *indivisible*. You cannot give half of a continued life to Bob and keep the other half to yourself. The only way in which you can assign equal shares is by vaporizing yourself, thus letting Bob fall to his death. That way, both of you end up with a share of zero. But this maximally fair solution is unacceptably wasteful.¹² If it is possible to strike a balance between fairness and efficiency, 'continuing one's life' is too valuable a resource to sacrifice for maximal fairness.

Broome suggests that when we are distributing an indivisible good between candidates with equally strong claims, we should hold a lottery in which all candidates are given an equal chance of winning the good. This avoids wastefulness, and the requirement of fairness is met to a considerable degree (Broome, 1991, pp. 97-8; original emphasis):

“[If the good in question is indivisible and we don't want to let it go to waste], the candidates' claims cannot all be equally satisfied, because some candidates will get the good and others will not. So some unfairness is inevitable. But a sort of partial equality in satisfaction can be achieved. Each person can be given a sort of surrogate satisfaction. By holding a lottery, each can be given an equal *chance* of getting the good.”

Following Broome, we may thus conclude that as a potential victim of a non-responsible threat, you should flip a coin. You—a morally responsible agent—are charged with the distribution of an indivisible good. By 'charged' I mean: It just so happens that an important good needs distribution, and that you will end up distributing it, in that your action or inaction will

¹²Note that it is not Pareto efficient.

inevitably determine how the good is distributed. Given this, you should aim at a distributively just solution, and should thus give an equal chance of getting the good to Bob and yourself. If you do this, your action properly reflects the equal moral standing Bob and you enjoy as individuals (cf. Broome, 1984, p. 38; Levine, 1984, p. 70; Davis, 1984, pp. 919-2; Lazar, 2009, p. 715) as well as the great value of the good in question (because you refuse to let it go to waste). If you shoot Bob without flipping a coin first, you rob Bob of his rightful chance to continuing his life. This chance is valuable: it represents the ‘surrogate satisfaction’ of Bob’s claim to the good of a continued life, given the scarcity and the indivisibility of the good in question. It follows that you must not shoot Bob without holding a lottery first.¹³

5 Comparing Innocent Threat Cases

So far, I have argued for the following claims:

1. Compared to an innocent bystander, Bob non-responsibly depends on you for his life. This weakens his moral standing vis-à-vis you.
2. Compared to an innocent bystander, Bob non-responsibly threatens your life. This weakens his moral standing vis-à-vis you further.

Here, I want to test whether our intuitions are in line with these claims. To do this, I will alter the Original Case presented in the introduction by removing the factors that I have argued work to Bob’s disadvantage. If (1) and (2) are correct, we should judge that Bob’s moral position is strengthened as we move from one case to the next.

¹³May you let yourself be crushed without flipping a coin first? It seems to me that while such a heroic act of self-sacrifice is not morally required, there is nothing wrong with your freely choosing it. Thomson (2008) would disagree. She argues that “a willingness to give up one’s life *simply* on learning that [someone else] will live if and only if one dies is a sign of a serious moral defect in a person” (ibid., p. 366, original emphasis), as it is indicative of the fact that one values one’s own life insufficiently. While I agree that proper respect for persons has to include proper self-regard, it does not seem to me that such self-regard is incompatible with a willingness to sacrifice oneself for an unspecified stranger. In the case at hand and given your innocence, I would think proper self-regard incompatible with a *denial* that you have a *rightful claim* to continuing your life, but not with the fact that you may *rightfully choose* to waive such a claim.

Case A

(α) *Bob does not profit from you;*

(β) *Bob is a threat to you*

The setup here is the same as in the Original Case, except for the bottom of the well now being laid out with protective padding. If you do nothing and let Bob crush you, you still die, but it is now the protective padding that cushions Bob's fall.

Given the setup of Case A, Bob is still intimately tied up with what poses a threat to you, but now your presence is no longer necessary for his survival. Rather, your absence would now guarantee his safety. If given a choice, Bob would thus want you gone. Without you there, Bob would not have to deal with the prospect of crushing (or else being vaporized by) an innocent stranger, and the protective padding would safely cushion his fall.

Intuitively, I feel that Bob's moral position vis-à-vis you is strengthened as we move from the Original Case to Case A. This means that my intuitions are in line with the idea that Bob's depending on you works to his moral disadvantage.

In addition to the arguments already presented, there is one further theoretical consideration that supports this intuitive judgement. Suppose we follow Warren Quinn (1989) and (more recently) Jonathan Quong (2009) in spelling out the Kantian idea of 'using someone as a mere means' in a counterfactual way, so that *X uses Y as a mere means whenever X harms Y without Y's consent, thereby achieving something X would not have been able to achieve in Y's absence*, Bob uses you as a mere means in the Original Case, but not in A. Under such a counterfactual definition of the term, my intuitive judgements thus suggest that in the Original Case, Bob's moral status is weakened by the fact that he uses you as mere means.¹⁴

¹⁴Quinn (1989, p. 344) refers to using someone in this way as "direct opportunistic agency". and claims that it is especially hard to justify. Quong (2009, p. 525) agrees that there is something "particularly wrong" with harming someone in this way. When it comes to profiting non-responsible threats, there is the complication that the threat lacks control over his action, and can thus not be accused of "direct opportunistic agency". But if a threat *harms his victim without the victim's consent, thereby securing his survival (that he would not otherwise have been able to secure)*, this seems to me sufficient to conclude that he—entirely non-responsibly—uses his victim as a mere means in a counterfactual sense. As I have argued in section 4, our moral standing vis-à-vis each other does not depend exclusively on what we do to each other as morally responsible agents; the effects we have on others merely in virtue of our existence as embodied creatures can be morally

In Case A, would it be wrong of you to flip a coin to determine who should live? As discussed, flipping a coin seems to me very much an upper limit of what could be required of Bob if he were able to control his fall; accordingly, a 50% chance of death seems to me very much an upper limit of what it would be permissible to impose on Bob in Case A. If you saw an option to give Bob a slight advantage over yourself, I think that you ought to go with that option instead of flipping a coin.

Case B

(α) Bob does not profit from you;

(β) Bob is not a threat.

This case is like Case A, except that the thugs throw a spear down the well before they throw in Bob. If you do not shoot your ray gun, your heart will be pierced by the spear, and Bob will then land on your dead body. You can either vaporize Bob and the spear at one go, or you can let your heart be pierced by the spear. The spear poses no threat to Bob. Moreover, if there were no spear, the impact of Bob's body on yours would cause you no harm.

In Case B, Bob and what's threatening you are entirely separable. That is, we can remove Bob from the scene without removing the threat, and we can remove the threat without removing Bob (so that Bob's presence is no longer sufficient for your life to be at risk). In Case B, if anything can be said to differentiate Bob from an innocent bystander, it seems to be mere spatial location—Bob is falling towards you, whereas innocent bystanders are usually assumed to stand some distance apart.

I feel that the Bob of Case B is morally indistinguishable from an innocent bystander whom you mustn't ride roughshod over. Intuitively, such a bystander has to incur moderate costs to move out of the way so that you can save yourself, but his duties are much more limited than those of a non-responsible threat. In B, it would therefore be impermissible to flip a coin to decide who of you should live.

significant also.

6 A Prerogative to Save One's Own Life?

I shall now consider an important objection to the idea that when faced with a profiting non-responsible threat, one ought to flip a coin. The objection can be put as follows: “From an impartial point of view, the claims of a non-responsible threat and his potential victim may indeed tie. But it does not follow from this that the potential victim has to treat both claims equally, as he is not required to take up an impartial perspective. Given that his own life is at stake, the victim may give more weight to his own interests, and may permissibly treat his own claim as the more significant one.” Nancy Davis (1984) and Jonathan Quong (2009) have both presented such a prerogative-based argument in favour of the permissibility of killing non-responsible threats in self-defence.¹⁵ In the following, I focus on Quong’s argument, and explain why I think it fails.

Quong (2009, p. 518) comes up with and defends a general principle called the *Principle of Defensive Killing* (PDK). The PDK states that “[you] can permissibly kill [another person] X if X will otherwise kill you”, provided that “(i) killing X is the only reasonable course of action that can save your life, and (ii) you have not waived or forfeited your permission to act in self-defense.” (ibid.). According to the PDK, you are allowed to vaporize Bob because unless you do so, Bob will kill you, and because your vaporizing him is in compliance with (i) a proportionality and (ii) a ‘no forfeiture’ clause.

Like Thomson, Kamm, and Tadros, Quong argues unambiguously in favour of the permissibility of killing a non-responsible threat. His approach nevertheless differs from theirs in at least two important respects. Firstly, while Thomson, Kamm, and Tadros claim that you may kill Bob because he threatens to *violate your rights* or because he is *liable to defensive harm*, Quong thinks you may kill Bob *simply* because he will kill you unless you kill him. Secondly, Quong (2009, p. 519) argues (contra Thomson and Tadros) that responsible agency is necessary for violating another’s rights as well as for becoming liable to defensive harm. To Quong, when Bob starts falling towards you, he neither threatens to violate your rights, nor does he become liable to defensive harm. Unlike Thomson and Kamm, Quong is therefore not committed to the counterintuitive claim that ‘ought’ need not imply ‘can’. Moreover, unlike Thomson, Kamm, and Tadros, he is not committed

¹⁵Levine (1984, esp. p. 74) presents a similar argument, but rejects it.

to the revisionist claim that responsible agency is not a necessary condition for liability.

Instead, Quong appeals to the idea of a prerogative to justify the PDK. When an agent's own life is at stake, Quong (2009, pp. 516-7) thinks that the "agent-relative value each person's life has for them" grounds a "powerful agent-relative permission to avoid sacrificing or significantly risking their own life for the sake of others [...]." While Quong does not argue that such a prerogative implies a general permission to kill an innocent person whenever doing so would serve to save one's own life, Quong does believe that an agent may permissibly kill an innocent person to save his own life just in case he does not use that person as a mere means. Whether an agent uses another person as a mere means is decided through consideration of a counterfactual: if the agent would not be able to secure his own survival if the other person were not present, the agent uses that person as a mere means. Quong (2009, p. 525) appeals to a Kantian idea of respect to argue that there is something "particularly wrong" with harmfully using someone as a mere means, so that our otherwise powerful prerogative to save our own life cannot extend to killings that involve such use. Quong then claims that according to his definition of using as a mere means, no one is used as a mere means in the cases that the PDK captures, simply because where X is what threatens to kill you, removing X does not make it impossible for you to secure your survival—rather, it guarantees it.

At first sight, Quong does not seem to miss the fundamental symmetry that I have argued characterizes non-responsible threat cases. Quong would agree that from an impartial point of view, Bob's and your claim to the good of a continued life tie. In fact, Quong insists that neither the non-responsible threat nor his potential victim lose their respective rights not to be killed, and that both threat and victim retain their right to self-defence, in that according to the PDK, they may permissibly infringe the other's right not to be killed (Quong, 2009, p. 525; see also Davis, 1984, p. 193). In sum, Quong seems to offer a clear and elegant solution to the problem of self-defence against non-responsible threats that is consistent with the view that the situation between a profiting non-responsible threat and his victim is roughly symmetric from an impartial point of view.

On closer inspection, however, our Original Case points to a weakness in Quong's account. Succinctly put, the problem is that Quong's PDK does

not match up with the rationale he provides in its favour. According to the PDK's *wording*, the situation between Bob and you is indeed symmetric as I have sketched it in the above paragraph. But according to its *rationale*, the situation is asymmetrical: given that Bob would be using you as a mere means if he were to kill you, he would—according to the PDK's rationale—have no right to kill you. If he could smash your ray gun, he would thus be disallowed from doing so, even if he knew that you will shoot him if he doesn't. The letter of the PDK thus allows certain killings that its justification forbids.¹⁶

To make matters worse, there are killings that seem intuitively impermissible, but that according to the PDK's justification we ought to endorse. Consider again Paul, the innocent bystander. If you vaporize the piano, you save your life without using Paul as a mere means, as Paul's presence makes no difference to your ability to vaporize the piano (Otsuka, 1994, p. 77; Hanna, 2012). In fact, you would welcome Paul's absence, as you could then vaporize the piano without also killing an innocent person. It follows that according to the PDK's rationale, it is permissible to vaporize Paul. I agree with most other moral philosophers that this is an unacceptable implication.¹⁷

Why does the PDK run into these problems? The underlying issue seems to me the following. Quong thinks that the agent-relative value each person's life has for them makes it sometimes permissible to kill an innocent person in order to save one's own life. He tries to substantiate what he means by 'sometimes'. both by formulating a rule-like principle (the PDK), and by providing a rationale in its favour. While it may be possible to come up with a rule-like principle that matches our intuitions about specific cases, the search for an adequate rationale with genuine explanatory power may well be a futile one. Killing an innocent person *simply because* one puts one's own interests first sounds like a gravely disrespectful thing to do, and I doubt that there is a convincing rationale why doing so should sometimes be thought morally justifiable (as opposed to merely excusable). For any

¹⁶This assumes that intentional agency is not a necessary condition for 'using as a mere means'. Quong might want to disagree with this. See footnote 14.

¹⁷Quong is aware that the PDK's rationale sanctions killing Paul, and is willing to accept this implication of the PDK (that is, he thinks that we ought to reconsider our intuitions about the permissibility of killing innocent bystanders such as Paul). See Quong, 2009, pp. 529-30.

proposed rationale, I thus conjecture that it will be possible to construct counterexamples to it, i.e. cases in which the rationale permits killings that we find intuitively wrong.

I therefore propose a different approach. I suggest we take as a starting point cases where substantial harm needs distribution among innocent individuals and where one may clearly put one's own interests first, to see if a plausible constraint on agent-relative prerogatives emerges. What are cases where substantial harm needs distribution among equally undeserving individuals, and where we do not doubt that the special value each person's life has for them grounds a prerogative to put one's own interests first? Suppose that you are out hiking in the woods with an innocent stranger when you chance upon a hungry bear (Boorse and Sorensen, 1988, p. 115). The bear starts charging, and it is evident that he will eat either you or the innocent stranger depending on who manages to run faster. Do the two of you have to flip a coin in order to determine who may run and who has to stay put? Of course not. You may simply run as fast as you can. It is not morally required that either of you be given a fair chance to survive; even if you are confident that you are able to run much faster than the stranger, you may still run for your life, thus leaving the other behind (Boorse and Sorensen, 1988, pp. 115-6). What it intuitively seems you must not do is stab the stranger with your hunting knife in order to slow him down.

To return to our Original Case, suppose the bottom of the well were large enough for you to step aside, so that you could save your own life without killing Bob, but by merely letting him die. In such circumstances, it seems to me again obvious that you would be permitted to secure your own survival, despite the fact that Bob is innocent, and despite the fact that stepping aside would mean his certain death. What makes the Original Case tricky is the fact that you cannot save your own life without *killing* Bob.¹⁸ Can these intuitions be said to capture a moral principle? At the very least, they are consistent with the following constraint on agent-relative prerogatives (CAP):

¹⁸Elsewhere, Quong agrees that the distinction between harming an innocent other and allowing him to be harmed is morally significant. See Quong, 2012, p. 63, fn. 29.

CAP. It is impermissible to depart from an impartial point of view and give more weight to your own interests if you can do so only by inflicting harm on another person that this person would not be duty-bound to suffer for your sake.

To violate the CAP is to make an exception of yourself in a morally objectionable way: it is to pursue your own interests in precisely the way you are morally asked not to pursue them. If you inflict harm on another person that this person is not required to suffer for you, you trespass on another person's sovereign territory that morality clearly flags as 'off-limits' for you. But why is each of us morally entitled to territory of their own? Why is our morality based on such a strong presumption against harming innocent others? Note first that we are social beings. As such, we lead our lives surrounded by other people. Second, we are rational creatures. As such, we have our own plans and projects, and we aspire to lead our lives in accordance with them. Third, we are embodied creatures: if someone harms our bodies, they harm us. Stringent rights to non-interference enable us to live together peacefully while still pursuing our own interests. Only if there is a general—as well as generally accepted—rule that we must avoid harming each other can we go about our lives without constantly seeing each other as potential threats.

If the CAP did not exist, or if people constantly ignored it, this would undermine the function of our stringent rights to non-interference. To violate the CAP is thus to pursue your own good in a way that is not universalizable. The situation presents itself differently if you are able to pursue your own good without trespassing on someone else's sovereign territory. In that case, a permission to put your own interests first need not undermine our general ability to pursue our own interests while living together peacefully. It follows that if you can pursue your own interests without trespassing on anyone else's sovereign territory, then you are in luck: an agent-relative prerogative may well apply to you.

Suppose this is roughly right, and something like the CAP does in fact apply. There are then no circumstances under which you have a prerogative to kill an innocent other in order to save your own life. Also, the only reason why you don't have to let Bob crush you, but may instead flip a coin to determine who of you should live, is because flipping a coin is the impartially right thing to do. When it comes to Bob and you, we can take up whatever perspective we want—the situation between the two of you is

always a toss-up.¹⁹

7 Conclusion

Given the abstract nature of the cases discussed in this paper, it may not be immediately obvious why clarifying the moral status of a profiting non-responsible threat is relevant to political philosophy. Bob is not only a peculiar creature from a moral point of view; he is also a rare occurrence in real life. But once we keep in mind that Bob is a stylized version of an innocent person who happens to impose on innocent others, it should become clear that the issues discussed in this paper are relevant to problems that can arise in any violent context, but that are maybe especially prone to arise in the context of terrorism.

For a fairly realistic example, suppose that terrorists have set a bus of civilians on a path to a steep drop. On an observation deck above the drop, a crowd of civilians is enjoying the view. If the bus crashes into the crowd, it will slow down sufficiently so that it doesn't go over the edge and into the ravine. If the bus crashes into the crowd, most of the passengers on the bus will get away with minor injuries, but many individuals in the crowd will be killed. If the bus goes over the edge, it is unlikely that any of its passengers will survive. Roughly the same number of people will be killed if the bus plows into the crowd as would be killed if it went over the edge instead. If the only way to stop the bus before it reaches the unsuspecting crowd is to blow it up (thus killing most of the passengers on the bus), should the bus be stopped, or should the decision be made to let it run into the crowd?

As a general rule, a government must not order the killing of a group of civilians in order to keep another group from being killed. In this paper, I have argued that in the just sketched example, we are dealing with an

¹⁹Given our morality's strong presumption against harming innocent others, wouldn't it be better for an otherwise uninvolved bystander not to intervene? I don't think so. A strong presumption against harming innocent others means that innocent people will not usually be duty-bound to bear much harm. But when in exceptional circumstances they *are* bound to bear a certain amount of harm, we should intervene to ensure that the right people are harmed to the right extent. Also, even if flipping a coin is the impartially right thing to do, an uninvolved bystander may still have an agent-relative prerogative not to intervene. That is, our morality may not be as demanding as to require that a bystander run the risk of having to kill an innocent person in order to give another innocent person a fair chance to survive.

exception to this general rule. The passengers on the bus are profiting non-responsible threats.²⁰ As such, they would have to suffer substantial harm to divert the bus away from the crowd if they saw an opportunity to do so. This grounds a justice-based reason to make them share in the harm that they are about to impose on the crowd. The fact that they depend on the crowd's presence for their lives weakens their moral status vis-à-vis the individuals in the crowd further. Given this, the Special Forces team on location should be instructed to give an equal chance of survival to both the civilians on the bus and those in the crowd.

What if the people in the crowd had the means to blow up the bus without outside help? According to the argument that I have presented in this paper, we have no prerogative to kill innocent others to save our own lives. But if the individuals in the crowd found a way to move out of the bus' way, they would be permitted to do that. I have suggested that when we can put our own interests first without inflicting harm on innocent others, we may well have a prerogative to do so.

²⁰In this example, it is not the passengers' bodies that endanger the people on the observation deck; it is the out-of-control bus that they are sitting in. Does it still make sense to characterize the passengers as threats? I think so. Qua moral agent, Bob happens to be trapped inside his out-of-control body, whereas the passengers happen to be trapped inside an out-of-control bus. Like Bob, who would have to steer himself away from you if he could, the passengers on the bus would have to steer the bus away from the crowd if they could. If Bob chose not to redirect himself at moderate cost to himself, he would be violating your rights to non-interference; if the passengers on the bus chose not to steer themselves away from the crowd at moderate cost to themselves, they would be violating the rights to non-interference of the people on the observation deck. This suggests that a non-responsible threat can be characterized as someone who is inside an object that he has non-responsibly lost control over and that now threatens to impose harm on innocent others.

References

- Boorse, C. and Sorensen, R. A. (1988). Ducking Harm. *The Journal of Philosophy*, 85(3): 115–134.
- Broome, J. (1984). Selecting People Randomly. *Ethics*, 95(1): 38–55.
- Broome, J. (1991). Fairness. *Proceedings of the Aristotelian Society, New Series*, 91: 87–101.
- Davis, N. (1984). Abortion and Self-Defense. *Philosophy & Public Affairs*, 13(3): 175–207.
- Draper, K. (1993). Fairness and Self-Defense. *Social Theory and Practice*, 19(1): 73–92.
- Hanna, J. (2012). The Moral Status of Nonresponsible Threats. *Journal of Applied Philosophy*, 29(1): 19–32.
- Kamm, F. (1987). The Insanity Defense, Innocent Threats, and Limited Alternatives. *Criminal Justice Ethics*, 6(1): 61–76.
- Kamm, F. (1992). *Creation and Abortion: An Essay in Moral and Legal Philosophy*. Oxford University Press, New York NY.
- Kamm, F. (1996). *Morality, Mortality Volume II: Rights, Duties, and Status*. Oxford University Press, New York NY.
- Kamm, F. (2012). *The Moral Target: Aiming at Right Conduct in War and Other Conflicts*. Oxford University Press, New York NY.
- Lazar, S. (2009). Responsibility, Risk, and Killing in Self-Defense. *Ethics*, 119(4): 699–728.
- Levine, S. (1984). The Moral Permissibility of Killing a “Material Aggressor” in Self-Defense. *Philosophical Studies*, 45(1): 69–78.
- McMahan, J. (1994). Self-Defense and the Problem of the Innocent Attacker. *Ethics*, 104(2): 252–90.
- McMahan, J. (2002). *The Ethics of Killing. Problems at the Margins of Life*. Oxford University Press, New York NY.

- McMahan, J. (2005). The Basis of Moral Liability to Defensive Killing. *Philosophical Issues*, 15(1): 386–405.
- McMahan, J. (2009). *Killing in War*. Clarendon Press, Oxford.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Basic Books, New York NY.
- Otsuka, M. (1994). Killing the Innocent in Self-Defense. *Philosophy and Public Affairs*, 23(1): 74–94.
- Øverland, G. (2011). Dividing Harm. *Journal of Moral Philosophy*, 8: 547–566.
- Quinn, W. S. (1989). Actions, Intentions, and Consequences: The Doctrine of Double Effect. *Philosophy and Public Affairs*, 18(4): 334–351.
- Quong, J. (2009). Killing in Self-Defense. *Ethics*, 119(3): 507–537.
- Quong, J. (2012). Liability to Defensive Harm. *Philosophy & Public Affairs*, 40(1): 45–77.
- Tadros, V. (2011). *The Ends of Harm. The Moral Foundations of the Criminal Law*. Oxford University Press, Oxford.
- Thomson, J. J. (1971). A Defense of Abortion. *Philosophy & Public Affairs*, 1(1): 47–66.
- Thomson, J. J. (1990). *The Realm of Rights*. Harvard University Press, Cambridge MA.
- Thomson, J. J. (1991). Self-Defense. *Philosophy and Public Affairs*, 20(4): 283–310.
- Thomson, J. J. (2008). Turning the Trolley. *Philosophy and Public Affairs*, 36(4): 359–374.