

Robust Rights and Harmless Wronging

Joseph Bowen

Stockholm University

This is a pre-publication draft of a paper forthcoming in *Oxford Studies in Normative Ethics*.

Abstract. This chapter examines a range of cases in which it appears one's rights against harm are violated by another's behaviour, even though this behaviour has done one no harm. These cases raise a serious problem for most theories of rights, though the problem is most pronounced on the Interest Theory of Rights. According to that theory, rights necessarily protect their holder's wellbeing. At first glance, one might think that the person's wellbeing cannot be said to be protected by the right in cases of harmless wrongdoing because they are not harmed in such cases—so, the necessary condition set for the ascription of a right is not satisfied. This paper puts forward a novel solution to this problem, the Safety Condition. The Safety Condition looks beyond what happens in the actual world to close worlds in order that individuals' wellbeing is robustly protected across circumstances that could just have easily come about.

Keywords. Rights; Harmless Wronging; Harm; Wellbeing; Safety; Interest Theory; Pre-emption; Pure Risk.

1. Introduction

Rights are important. Part of why rights have this importance is that they protect us from harm—they protect our interests. This chapter examines a range of cases in which it appears one's rights against harm have been violated by another's behaviour, even though this behaviour has done no harm. Call these cases of “harmless wrongdoing”.

Plane Crash (Preempted Harm). Passenger is about to board a plane. Attendant takes a disliking to Passenger, so denies her admittance onto the plane. On departure, the plane crashes and everybody on board dies.

Roulette (Pure Risk Imposition). Target is asleep. Her housemate, Shooter, comes into her room and decides to play Russian roulette with her. Luckily, no bullet is fired. Shooter, content with having had a round of roulette, will never play roulette again.

Intuitively, Passenger and Target have their rights violated.¹ However, Passenger is better off in the world in which her rights are violated than that in which they are not. Target's life is as it would have been had Shooter not made Target the subject of her risky behaviour.² Given the standard *Counterfactual Account of Harm*, on which *Y* harms *X* iff *Y* makes *X* worse off than *X* would have been had *Y* not acted as she did, Passenger and Target are not harmed by the violation of their rights.

Cases of harmless wronging raise a serious problem for most theories of rights (namely, those theories that have wellbeing play some grounding role in rights). For example, most theories say that, other things being equal, the stringency of a right corresponds to the harm that would befall its holder were that right to be violated (Thomson 1990, 149–75; Kamm 2007, 249–75). Cases of harmless wronging wreak havoc with the intuitively plausible stringencies of rights. What is the stringency of Passenger's and Target's rights against Attendant's and Shooter's behaviour?

¹ An anonymous reader for this volume would like to hear more about why it is intuitive that there are right violations in *Plane Crash* and *Roulette*. Here is another way at getting at the intuition. Many hold that one is wronged by some action only if one holds a right against that action. I imagine Passenger would not think, merely, that Attendant acted wrongly, nor do I believe Target would think Shooter merely acted wrongly (were she to find out what happened). Rather, they would feel *wronged*. This suggests, assuming the stipulated connection between rights and wrongs, Passenger and Target have their rights violated.

² Some think that risk of harm is *itself* harmful, e.g., (Finkelstein 2003). Elsewhere, I argue this is mistaken (Bowen Forthcoming).

The problem caused by cases of harmless wrongdoing is most pronounced on the dominant theory of the *nature* of rights, the Interest Theory of Rights. The problem is most pronounced for the Interest Theory because it is the theory that most depends on the intuitive connection between rights and harm that cases of harmless wrongdoing call into question.

Interest Theory (Canonical). X has a right against Y that $Y \Phi$, only if (and because) X 's wellbeing (her interests) is of sufficient weight to place Y under a duty to Φ .³

The Interest Theory has a lot going for it. As well as having an intuitively plausible account of the grounds of rights, it explains some fundamental structural features of rights and their correlative directed duties. First, it gives a good account of why Y *owes* her duty to X —the duty exists because of X . Second, it explains why, through infringing her duty, Y does not merely act wrongly but *wrongs* X — Y has failed to respond to morally salient duty-grounding features about X .

The Interest Theory runs into trouble with cases of harmless wrongdoing. Neither Passenger's nor Target's wellbeing is protected by their putative rights that are violated. Because of this, it is hard to see how either person's wellbeing is of sufficient weight to place anyone under a duty. So the necessary condition set for a right by the Interest Theory is not satisfied. Call this the Problem of Harmless Wronging for the Interest Theory. The problem is that of accommodating our intuitions that Passenger and Target have their rights violated, given

³ I omit that Y 's duty is *owed* to X as I see the Interest Theory as an account of *what it is* to owe a duty to another. Were we to allow "sufficient weight" to do a lot of heavy lifting, we could make Interest Theory (Canonical) necessary and sufficient for rights.

a commitment to the Interest Theory. More generally, the problem is that of accommodating agents having rights against harmless wrongs, given a commitment to the Interest Theory.

This chapter offers a principled solution to the Problem of Harmless Wronging by revising the Canonical statement of the theory with what I call the Safety Condition. The Safety Condition looks beyond what happens in the actual world to close possible worlds to normatively ensure that people's wellbeing is robustly protected across circumstances that could easily come about. In sections 2 and 3, I introduce the Safety Condition by showing how it deals with preempted harm and pure risk imposition. By the end of section 3, we see preemption and pure risk differ from mundane cases of rights against harm in symmetric ways. Section 4 begins to offer an account of why rights might be sensitive to modality in the way the Safety Condition prescribes, before defending the Safety Condition against two objections.

A preliminary: Suppose the Interest Theory fails. We should nonetheless want harm to explain Passenger's and Target's rights, even if we do not think *all* rights are grounded in their holder's wellbeing (and especially in *Hitmen*, introduced below). So those who deny the Interest Theory should still want a solution to the Problem of Harmless Wronging. Suitably refined, Safety Condition offers such an explanation.

2. Preemption

Let us begin by looking at the problem posed by our example of preempted harm, *Plane Crash*. An initial thought is that the problem posed by preemption is simply a problem with the Counterfactual Account of Harm. We need only refine or replace that. I think, though

do not argue here, that we have good reason to hold onto the Counterfactual Account.⁴ After introducing the Safety Condition, I return to this question of refining the Counterfactual Account.

Another tempting thought is that, although Passenger is not harmed all things considered by being denied admittance onto the plane, she is made worse off in a regard (she is harmed *pro tanto*). She suffers inconveniences that she would not have suffered were she not to have been arbitrarily denied admittance onto the plane.⁵ Although the extent to which she is made worse off is insufficient to leave her harmed all things considered, there is nonetheless *this regard* in which she is harmed. Perhaps she ought to be afforded a right protecting those aspects of her wellbeing. If we want to say this, we would need to endorse the following revision to the Interest Theory (Canonical):

Pro Tanto Thesis. X has a right that against Y that $Y \Phi$, only if (and *because*) some aspect of X 's wellbeing (an interest) is of sufficient weight to place Y under a duty to Φ .⁶

Because Passenger is made worse off in a regard, she can be afforded a right protecting those aspects of her wellbeing.

While the Pro Tanto Thesis might solve *Plane Crash* in that it accommodates the intuition that Passenger has her rights violated, it will not solve all cases of preempted harm. It provides a solution to *Plane Crash* only because Passenger is made worse off along some specific dimension—a specific dimension that she would not have been made worse off in

⁴ For discussion, see, among others, (Bradley 2012; Tadros 2016, 201–23).

⁵ We need to be careful not to introduce interests *created by* what we are owed, e.g., that Passenger has an interest in not being arbitrarily discriminated.

⁶ E.g., (Raz 1992, 129).

were Attendant to have allowed her onto the plane. However, this feature does not hold in all cases of preempted harm. Consider

Hitmen. Suppose that we have two hitmen. Hitman₂ admires Hitman₁. Hitman₂ secretly follows Hitman₁ on every job she has in the hope that, one day, Hitman₁ will fail to complete a hit and she will be able to do so instead, thereby impressing Hitman₁.

For any victim that Hitman₁ is contracted to kill (call her Victim), what aspect of her well-being is setback by Hitman₁ that would not have been setback by Hitman₂? I am sceptical there is any such aspect, assuming Hitman₂ would have completed the hit the instant Hitman₁ fails to. Accordingly, the Pro Tanto Thesis cannot account for why Victim has a right against Hitman₁ killing her. We ought to look beyond the Pro Tanto Thesis.⁷

Returning to *Plane Crash*, the problem of preemption arises because we focus on a comparison between the following two worlds:

World 1. Attendant denies Passenger admittance onto the plane. The plane takes off and crashes, killing everybody onboard.

World 2. Attendant does not deny Passenger admittance onto the plane.

The plane takes off and crashes, killing everybody onboard.

⁷ Two notes. First, the way I am thinking about pro tanto harm is that we see if there is a regard that Victim is worse off in that she would not otherwise have been worse off in. Bradley thinks about pro tanto harm differently: ‘Something is a [pro tanto] harm for a person if and only if either (i) it is intrinsically bad for that person, or (ii) it brings about something intrinsically bad for that person, or (iii) it prevents something intrinsically good for that person’ (2009, 66). If we suppose *Hitmen* involves painless death, only condition (iii) is relevant for our purposes, since death is only extrinsically good or bad. Since Hitman₁’s shooting Victim does not *prevent* anything intrinsically good for Victim because Hitman₂ would have shot Victim anyway, condition (iii) is not satisfied. So, Victim is not harmed pro tanto by Hitman₁ shooting her on Bradley’s account, either.

Second, a reader suggests in *Plane Crash* Passenger is harmed by not being let onto the plane, though the short-term set back of her interests leads to a long-term improvement. But this analysis will not help with *Hitman*.

Passenger is no worse off in world 1 through Attendant's denying her admittance onto the plane than she would have been in world 2. However, world 2 is not the only possible world available for comparison with world 1. We might look to

World 3. Attendant does not deny Passenger admittance onto the plane.

The plane takes off and lands safely at its destination.

Passenger is worse off in world 1 than she is in world 3. The extent to which she is worse off is sufficient to place Attendant under a duty not to deny her admittance onto the plane. World 3 is what we might call the *normal* counterfactual world. Perhaps we might revise the Interest Theory (Canonical) to respond to this fiat of the case.

Normality Thesis. X has a right against Y that $Y \Phi$, only if (and *because*) X 's wellbeing (her interests) is, *under normal circumstances*, of sufficient weight to place Y under a duty to Φ .

While the Normality Thesis does accommodate our intuition that Passenger has a right against Attendant, it will not solve all cases of preempted harm. For, harm may be preempted in the normal world and yet we still want to say that people should be attributed rights against those preempted harms. Consider *Hitmen*. Because Hitman₂ *always* follows Hitman₁ on every job, any victim that Hitman₁ is contracted to kill will not have a right against Hitman₁'s action. Victim would be no worse off under normal circumstances were Hitman₁ not to shoot, since Hitman₂ would make Victim worse off to an equal extent. The harm, then, is preempted under normal circumstances.

One might reply that it is not obvious that the circumstances described in this case are normal of hitmen in general, but only of Hitman₁. Perhaps we ought to restrict the reference class to *normal* hitmen. But this will not do, either. Suppose that, so eager to make sure that their hits are completed, assassination agencies begin to *always* send their hitmen

out in pairs. This means that for any victim it is normal that, were the first hitman to fail, another hitman would always be there to kill that victim.⁸

A less fanciful example than *Hitmen* is to imagine a polluting factory. Suppose these kinds of factories always pop up near each other, for example, because they always require a river to dump waste in, cheap land, proximity to work forces, and so on. On the supposition that one polluting factory is sufficient to harm on its own, any harm caused by a particular factory would under normal circumstances be preempted by another factory. This is because there would always be another factory nearby, itself polluting, that would have caused an equal harm.⁹

Let us set out the problem with the Normality Thesis. We were faced with a comparison between two worlds:

Hitmen World 1. Hitman₁ fatally shoots Victim. (Hitman₂ was waiting in the wings.)

Hitmen World 2. Hitman₁ does not shoot Victim. Hitman₂ fatally shoots Victim.

Victim is no worse off in Hitmen world 1 than she is in Hitmen world 2. Hitmen world 2 is what would happen under normal circumstances were Hitman₁ not to have shot. However, as with *Plane Crash*, there is another close world that we can appeal to,

Hitmen World 3. Hitman₁ does not shoot Victim. Hitman₂ does not shoot Victim.

⁸ This holds on most accounts of normality, including non-statistical ones, e.g., (Smith 2016, 38–45).

⁹ “Cancer Alley” comes to mind, an 85 mile stretch of the Mississippi that produces one quarter of America’s petrochemicals.

Victim is worse off in Hitmen world 1 than she is in Hitmen world 3. The question is how we make reference to *these* close worlds (Plane Crash world 3 and Hitmen world 3) to accommodate the intuitive verdict that Passenger and Victim have rights against the relevant conduct.

I suggest we make an appeal to *modal safety*. The idea behind safety is nicely explained by Timothy Williamson:

Imagine a ball at the bottom of a hole, and another balanced on the tip of a cone. Both are in equilibrium, but the equilibrium is stable in the former case, unstable in the latter. A slight breath of wind would blow the second ball off; the first ball is harder to shift. The second ball is in danger of falling; the first ball is safe. Although neither ball did in fact fall, the second could easily have fallen; the first could not.

(Williamson 2000, 123)

There is a danger an event will occur if that event does occur in some sufficiently similar case. In much the same way as the ball is not safely balanced on the top of the cone, Passenger's and Victim's wellbeing is not safely protected—though they are not actually made worse off, there is a danger they could have been. And, it is plausible that rights ought to safely protect their wellbeing.

We can make appeal to safety by revising the Interest Theory (Canonical):

Interest Theory (Safety). X has a right against Y that $Y \Phi$, only if (and *because*) Y 's not Φ -ing causes X to be worse off than she would have been in at least one close world, and the difference in X 's wellbeing is of sufficient weight to place Y under a duty to Φ .

For brevity, I call this the Safety Condition. It works by comparing how X fares when Y does not act as the would-be duty dictates with how X would have fared in close worlds in which Y acts as the would-be duty dictates. Call the world in which Y does not act as the would-be duty dictates our *world of evaluation*. Call the world in which Y does act as the duty

dictates our *world of comparison*. (This jargon becomes useful in the following section.) In *Plane Crash* and *Hitmen*, the world of evaluation is world 1, the actual world. We evaluate how our potential right-holder fares in this world through comparison with close worlds in which the potential duty-bearer acts as the would-be duty dictates (hence the names: world of evaluation and of comparison). Since Passenger and Victim are worse off in world 1 than they are in world 3, and since the extent to which they are worse off is of sufficient weight to place Attendant and Hitman₁ under their respective duties, Passenger and Victim hold rights. They hold these rights in the actual world.

As with the Interest Theory (Canonical), the Safety Condition is only a necessary condition.¹⁰ So, it does not follow that Passenger and Victim hold rights just because the Safety Condition is satisfied. But we can assume the other conditions necessary and jointly sufficient for rights are satisfied. We turn to these other conditions in section 4.3.

Above, I questioned whether we ought not refine the Counterfactual Account of Harm. One might wonder whether we ought to endorse the Safety Condition as an account of harm rather than as a refinement to the Interest Theory.¹¹ Extensionally, there is nothing between building safety into our account of rights or our account of harm. I would not be too worried if one takes all I say about the Safety Condition and builds it into their account of harm. However, let me offer two reasons why I prefer building it into our account of rights.

First, I do not believe that merely being subjected to risk is itself harmful (Bowen Forthcoming). As we see in the following section, the Safety Condition is satisfied in cases of pure risk, such as *Roulette*. If we build the Safety Condition into our account of harm, this

¹⁰ Again, we could allow “sufficiently weighty” to do some heavy lifting, and make the Safety Condition necessary and sufficient.

¹¹ This is something like Tadros’ *Complex Counterfactual View of Harm* (2016, 177).

means that merely being subjected to risk is itself harmful. So, it is preferable to build the Safety Condition into our account of rights—and say, though Shooter does not harm Target, she does violate her rights—rather than our account of harm.

Second, as we also see below, it appears that the Safety Condition overgenerates. If we build safety into our account of harm (and benefit), it overgenerates harms (and benefits). If we build safety into our account of rights, it overgenerates rights. I argue we can explain away these overgenerations. However, the sorts of considerations I appeal to in order to explain away these overgenerated rights are not the sort of features that ought to affect whether or not one has *harmed* another.¹² So, it is better to build the Safety Condition into our account of rights.

3. Risk

We have seen how the Safety Condition deals with cases of preempted harm. In this section, we see how the Safety Condition deals with cases of pure risk imposition.

In *Roulette*, Shooter comes into Target's room and plays a round of Russian roulette with her. Since no bullet is fired, Target is no worse off than she would have been had Shooter not subjected her to that risk. Because of this, Target's wellbeing is not protected by her right not to be subjected to gratuitous risk of harm. Given the Interest Theory (Canonical), Target has no right against Shooter that Shooter not play roulette with her.

One might hold Target has a right against Shooter because, given the best available evidence to Shooter, in expected terms Target's wellbeing is of sufficient weight to place her

¹² An example, preempting subsection §4.3: while we may think one's intentions or knowledge is relevant to whether others hold rights against one, it ought not affect whether one *harms* others.

under a duty not to play Russian roulette with Target. On this view, rights are determined from the evidence-relative perspective, rather than the fact-relative perspective as I have been assuming.¹³ I think there are problems with Evidence-Relative Views. For example, these views imply what rights I hold *change* as new evidence-becomes available to others, rather than others getting better evidence about what rights exist. These views also imply, if the best available evidence does not support that someone's behaviour will harm me (or, for other sorts of reasons, does not support that I have a right against being harmed), I have no right against that harmful behaviour. And there are lots of cases in which this looks very implausible (Bowen 2021). I do not say more about Evidence-Relative Views here. Instead, notice that even those very confident in Evidence-Relative Views need to see how it compares to the best alternative views. We can see the Safety Condition as offering an alternative fact-relative view of rights.

The Safety Condition can accommodate this verdict that Target has a right against Shooter that Shooter not subject her to such risky behaviour. Showing this requires looking back to preemption. Like Passenger in *Plane Crash*, Target is no worse off than she would have been had the duty correlative to her right been respected. But, unlike in cases of preemption, in which another event causes her to be at least as badly off as she would have been had the duty not been respected, this is because the risked harm does not materialise. We are comparing, then:

Roulette World 1. Shooter plays Russian Roulette. No bullet is fired.

Roulette World 3. Shooter does not play Russian Roulette.

¹³ E.g., (Zimmerman 2014; Quong 2015; van der Vossen 2016; Quong 2020). Quong's view is actually something of a hybrid. For discussion, see my (2021).

Target is no worse off in Roulette world 1 than she is in Roulette world 3. However, as with preemption, there is an extremely close world against which Target is made worse off because of Shooter playing roulette with her—namely, the world in which there *is* a bullet in the chamber when Shooter pulls the trigger:

Roulette World 2. Shooter plays Russian Roulette. Shooter fatally shoots Target.

Unlike when dealing with preemption, in which we compare world 1 with world 3, in cases of pure risk imposition we compare world 2 with world 3. Target is worse off in Roulette world 2 than she is in Roulette world 3. The extent to which she is worse off is sufficient to place Shooter under a duty not to make Target the subject of her risky behaviour. The Safety Condition can accommodate this verdict. It requires recognising Roulette world 2 as our focus when \mathcal{I} fails to act as the duty dictates instead of the actual world, Roulette world 1—it requires recognising Roulette world 2 as the *world of evaluation* rather than the actual world, Roulette world 1.

So, in cases of preemption, our focus when \mathcal{I} does not act as the duty dictates (our world of evaluation) is world 1, the actual world. In cases of pure risk imposition, our focus is world 2 (a world close to the actual world). It matches the Safety Condition for our focus on what happens when \mathcal{I} fails to act as the duty dictates to include worlds close to the actual world, rather than only the actual world itself. This is because we could easily have been in those close worlds—it could easily have been that there was a bullet in the chamber when Shooter pulls the trigger. While different events occur between Roulette world 1 and world 2 (Victim is not and is shot), Shooter *acts*, in the relevant sense, in the same way in both worlds. And this is in keeping with the Safety Condition: we are unsafe to the extent that our wellbeing is not safely protected; as Shooter could easily have shot Target, Target's wellbeing was not safely protected; by allowing us to focus on worlds in which Shooter acts

in the same way close to the actual world, rather than only the actual world, Target's wellbeing is safely protected. So far so good.

At this stage, one might wonder how we determine which world to focus on as our world of evaluation when \mathcal{I} does not act as the duty dictates. Why focus on the actual world in cases of preemption but focus on worlds close to the actual world in cases of pure risk? But we do not need to determine which of these worlds to uniquely focus on. While the Safety Condition speaks of *the* world of evaluation, this need not imply there is only one world of evaluation. Rather, there are multiple worlds of evaluation—those close worlds in which the potential duty-bearer performs the act that she may be under a duty not to perform. (This includes the actual world, for the actual world is close to itself.)¹⁴ In *Roulette*, worlds of evaluation are close worlds in which Shooter puts a bullet in the cylinder of the gun, spins the cylinder, and pulls the trigger. In *Plane Crash*, worlds of evaluation are close worlds in which Attendant denies Passenger admittance onto the plane. In *Hitmen*, worlds of evaluation are close worlds in which Hitman₁ shoots Victim. Again, this is in line with the Safety Condition: it ensures Target's wellbeing is robustly protected. We do not need to determine which world to focus on when \mathcal{I} does not act as the duty dictates since the Safety Condition looks to *all* close worlds in which the duty-bearer acts in the same way.

With this in place, we can better appreciate how the Safety Condition works. It makes for a two-part comparison. We have worlds of evaluation and worlds of comparison. Worlds of evaluation are worlds in which the potential correlative duty-bearer performs the action that she may be under a duty not to perform. Worlds of comparison are worlds in which the potential correlative duty-bearer does *not* perform this action. Whereas the canonical statement of the Interest Theory focuses only on the closest worlds in which the duty

¹⁴ While this is true on most accounts of closeness, it is not necessarily true, e.g., (Smith 2016).

is and is not respected, the Safety Condition looks also to other worlds (sufficiently) close to the closest worlds in which the duty is and is not respected. And if the potential right-holder is worse off in a world of evaluation than she is in a world of comparison, and the extent to which she is worse off is sufficient to place the potential correlative duty-bearer under a duty, the Safety Condition is satisfied.

As a point of comparison, consider a mundane case in which one's rights against harm are violated: suppose Threatener punches Innocent. Here, we compare the actual world, in which Threatener punches Innocent, with the closest counterfactual world in which Threatener does not punch Innocent. Our world of evaluation is the closest world (to the actual world) in which Threatener punches Innocent: the actual world. Our world of comparison is the closest world (to the world in which Threatener punches Innocent) in which Threatener does not punch Innocent. Since Innocent is sufficiently worse off in the world in which Threatener punches her than in the world in which Threatener does not punch her, Threatener is under a duty not to punch Innocent correlative to Innocent having a right that Threatener not punch her.

In cases of preemption, we compare how our potential right-holder fares not with how she fares in the closest counterfactual world in which the potential duty is respected, but with a different counterfactual world. In cases of pure risk, we do use the closest counterfactual world in which the potential duty is respected. However, we do not use the actual world as our world of evaluation, but some world close to the actual world. Inasmuch, preemption and pure risk differ from the mundane case in symmetrical ways. They are cases of harmless wronging in symmetrical ways to each other. (We could also have a case of preempted pure risk.)

To be precise when expounding the Safety Condition, I have used some jargon. Here is a more natural gloss on the Safety Condition. For someone to hold a right against us that we do not perform some action, we look to whether our performing that action could easily leave them sufficiently worse off to place us under a duty.

4. In Defence of Safety

Whereas the standard version of the Interest Theory and refinements considered above fall foul of the Problem of Harmless Wronging, the Safety Condition correctly generates rights in our cases of preemption and pure risk imposition. Further, the Safety Condition offers us a unified account of *why* people are attributed rights across these different types of case: even if one is not made worse off as things turn out, one could easily have been made worse off. The Safety Condition's *principled* extensional accuracy is the primary virtue of the account that I would like to stress in this chapter.

In this section, I defend the Safety Condition against two objections: that it under- and overgenerates rights. Before that, I say a little about *why* modality might matter for rights in the way that the Safety Condition prescribes.

4.1 *Why Safety?*

One important feature of the Safety Condition is that it removes an objectionable form of luck from rights. If we assume the Interest Theory (Canonical) then, through sheer luck, Attendant does not violate Passenger's right not to be denied admittance onto the plane. It could easily have been that Attendant *did* harm Passenger, and so would have violated her rights. Similarly, through sheer luck Shooter does not violate Target's rights. It could easily have been that there was a bullet in the chamber when Shooter pulled the trigger, and so Shooter would have violated Target's rights. By focusing on more than what happens in the

actual world, the Safety Condition removes this objectionable form of luck from rights (and right-violations). This gives us good reason to endorse the modal character of the Safety Condition.

One might wonder why it matters that rights do not depend on (objectionable) forms of luck. Some hold that others are worse off to the extent that their wellbeing depends on luck (Oberdiek 2017; Lazar 2017). This is usually couched in terms of luck being antithetical to control, and control being necessary for one to lead an autonomous life. If we were to go this way, it would leave the Safety Condition unmotivated as concerns those for whom autonomy has little or no value (for example, very young children and those with severely damaged rational capacities) (Bowen Forthcoming). And, regardless, I believe that at least rights against *harm* should be grounded in facts about wellbeing in general, rather than autonomy.

Instead, I propose it matters that rights do not depend on luck in this objectionable way—or, that rights depend on modality as the Safety Condition prescribes—because this formally requires that we are sensitive to others’ wellbeing: that it is not merely that we do not harm others, but that we could not easily have harmed others.¹⁵

The Interest Theory starts with the idea that others’ wellbeing is sufficiently important to place others under duties. You need not be an Interest Theorist to believe that. The Interest Theorist is distinctive because they say, a nice explanation for when and why you owe someone a duty is that their wellbeing is the grounds of the duty. Commonly, it is taken that duties are just a special type of reason: they have something like exclusionary

¹⁵ In saying that duty-bearers are sensitive to others’ wellbeing, one might think I am confusing my modal notions. *Y*’s Φ -ing *safely* does not harm *X* iff there is no close world in which *Y*’s Φ -ing harms *X*. *Y* is *sensitive* to her Φ -ing not harming *X* iff, were Φ -ing harmful, *Y* would not Φ . That *Y* safely does not harm *X* does not mean that she is sensitive to her Φ -ing not harming *X*. To be more precise, I should say the Safety Condition makes *Y* sensitive to her Φ -ing not being harmful in sufficiently close worlds.

weight, they leave a moral remainder when not acted on, they are demandable, and so on. The Safety Condition's focus on what could easily happen makes duty-bearers' reasons of this kind more sensitive to others' wellbeing than is the case on the modally undemanding, canonical statement of the Interest Theory. Since the Interest Theory began with the idea that others' wellbeing is very important—it both places us under duties and exclusively makes those duties owed to others—it is in keeping with the Interest Theorists' motivations to endorse the Safety Condition.

Think back to Williamson's example from above. Suppose you have asked two people to put the ball somewhere and keep it still. When all else is equal, the person who puts the ball at the bottom of the hole has taken more care to ensure that the ball is still than the person who balanced the ball on the top of the cone. Analogously, the person who turns out not to harm others, but could easily have, such as Attendant, has taken less care not to harm others than those who robustly do not harm others. This offers an attractive picture of what we *owe* to others—that we take care not to leave others worse off through our actions than otherwise could have been the case.

This is just the beginning of an account of *why* rights respond to modality in the way the Safety Condition prescribes. But it is helpful to have more than extensional accuracy behind the Safety Condition before turning to our objections.

(One question this prompts: does it matter that we are sensitive to others' wellbeing only when we are uncertain of how our actions will turn out? When harm is preempted, I am tempted to say we have rights against others' actions even if they know their harming us is preempted. This explains why, even if someone's contribution to my being harmed is preempted, I could use them as a means to prevent myself from being harmed. I am less sure when it comes to pure risk. How much would I pay not to be subjected to a merely

modal risk?¹⁶ Little, if anything. This suggests worlds do not matter when we know they are not going to come about. On this way of thinking about things, while evidence affects which possible worlds are relevant to rights, it is very far from Evidence-Relative Views mentioned in section 3. On this version of the Safety Condition, rights depend only on the facts—it is just that modal facts about what could otherwise have been the case might lose their relevance if we *know* those worlds are merely modal.)

4.2 *Very Preempted Harm*

In *Plane Crash*, the Safety Condition works by comparing the actual world with the close world in which Passenger boards the plane and it does not crash. In *Hitmen*, the Safety Condition works by comparing the actual world in which Hitman₁ shoots Victim with the close world in which Hitman₁ does not shoot Victim and Hitman₂ does not shoot Victim. In both cases, we might say, we are looking past the *closest* (counterfactual) world to some other close (counterfactual) world in which the preempting harm does not occur. However, here lies the recipe for a counterexample: make the harm *very* preempted.

Smokin' Aces. Hitman₁ is contracted to kill Victim. Unbeknownst to Hitman₁, there are one hundred other hitmen waiting in the wings. Each is ready to kill Victim if the previous hitman fails.

In *Smokin' Aces*, the Safety Condition has to compare:

(*Smokin' World 1*) Hitman₁ shoots Victim (Hitmen₂₋₁₀₀ were waiting in the wings),

with:

(*Smokin' World 101*) Hitman₁-Hitman₁₀₀ all do not shoot Victim.

¹⁶ Thanks to Peter Graham for putting it like this.

Victim is worse off in Smokin' world 1 than she is in Smokin' world 101— Smokin' world 101 is the closest world that allows us to say this. However, one might hold that Smokin' world 101 is *not* a close enough world for the purposes of satisfying the Safety Condition. So, Victim will not have a right that Hitman₁ not shoot her.

Why is Smokin' world 101 is not a close world? Consider David Lewis's view of closeness. Begin by supposing that the worlds under consideration have deterministic laws: they start and play out following deterministic laws of nature. We discover the closeness of two worlds at some time, t , by reference to the number and size of the violations of those laws of nature that would be required at t to render those worlds convergent after t (Lewis 1979, 472). We need 100 miracles to get us from Smokin' world 101 to Smokin' world 1. And perhaps that is too many miracles.

One way to address this Problem of Very Preempted Harm is to lean on our theory of closeness to avoid the verdict that the world in which the harm is not preempted is not close. For example, because the differences between Smokin' world 1 through to Smokin' world 101 are *themselves* so similar, one could argue Smokin' world 101 is *not* actually all that far from Smokin' world 1. Call this the *Similarity Strategy*. Taking this strategy, we have not relaxed how close a world needs to be for it to be close enough to satisfy the Safety Condition (which would make the view more likely to overgenerate rights), but suggested that Smokin' world 101 is not actually that far away, so will satisfy the Safety Condition. We might even say, how far away these worlds are tends towards some limit no matter how many hitmen we add.

There is at least one limitation with the Similarity Strategy. It works in *Smokin' Aces* because the many different preempting acts are relevantly similar, thereby diminishing their impact on how far away worlds are. But maybe this is just a fiat of *Smokin' Aces*. Perhaps one

could devise a case in which all the different things that cause the harm to be preempted are of a different character, meaning that the Similarity Strategy will be of little help. We could suppose that were Hitman₁ not to have killed Victim, a boulder would have fallen on her and, had the boulder not fallen, lightning would have struck her, and so on.

Another way to resist the Problem of Very Preempted Harm is to revise the Safety Condition. The element that cases of very preempted harm put pressure on is that X must be worse off in at least one *close* world. We could revise this closeness element:

Interest Theory (Safety, Relevance Variant). X has a right against Y that $Y \Phi$, only if (and *because*) Y 's not Φ -ing causes X to be worse off than she would have been in at least one relevant world, and the difference in X 's wellbeing must be of sufficient weight to place Y under a duty to Φ .

Since there is no stipulation that the world of comparison needs to be a close world, we need not worry that Smokin' world 101 is far away from Smokin' world 1. It turns only on whether we think Smokin' world 101 is relevant. And it seems relevant.

There is an obvious worry with the Relevance Variant: how do we work out which worlds are relevant? Despite the Relevance Variant faring better with the Problem of Very Preempted Harm than the standard Safety Condition, I am not sure that there is a non-circular way to answer this question. So while the Relevance Variant may score well extensionally, it has less explanatory power than the Safety Condition.

But perhaps this worry is premature. Speaking of the safety condition on knowledge, Williamson says: 'In many cases, someone with no idea of what knowledge is would be unable to determine whether safety obtained [...] One may have to decide whether safety obtains by first deciding whether knowledge obtains, rather than vice versa' (2009, 305). On this way of understanding things, safety can be thought of as offering a circular account of

knowledge. Return now to the Relevance Variant of the Safety Condition. Even if we cannot know whether one is owed a duty and holds a correlative right by the lights of the Relevance Variant alone, that does not mean that it is ontologically circular. We can still maintain that whenever one is owed a duty, and holds a right correlative to that duty, the duty is grounded in how one fares across relevant worlds. This is so even if we need a prior understanding of whether one is owed that duty to know whether the world is as relevant.

There is a third response to Very Preempted Harm, one that does not require moving to the Relevance Variant: we might accept our problem's conclusion. When harm is preempted to a great degree, along many different dimensions, perhaps people do not have rights against that harm. Let me offer two remarks to make this bullet is easier to swallow.

First, above we considered only one feature of closeness (the Similarity Strategy) that helps solve the problem. We might hold out hope that there are other features of the semantics of closeness that will provide a solution to our problem.

Second, and more substantially, we are attempting to offer a *reductive* account of rights (and directed duties)—an account that explains rights and what it is to owe a duty to another person by appealing to some other feature(s). The Interest Theory (with the Counterfactual Account) began with the idea that rights are difference makers. The Safety Condition adds, “rights are difference makers or *could-easily-have-been* difference makers.” The Problem of Very Preempted Harm presses, “Well, what if the putative duty-bearer couldn't easily have been a difference maker?” Perhaps it should not be surprising that we need to accept some counterintuitiveness along this line. When someone really is doomed to suffer some

bad fate, perhaps they do not have rights against us that we not make that fate preempted to more of an extent.¹⁷

4.3 Overgenerations

Whereas the Problem of Very Preempted Harm suggests that the Safety Condition will undergenerate rights, the Safety Condition may also overgenerate rights. The Safety Condition requires only that there is one close world in which the right-holder fares sufficiently better through the duty-bearer acting as the duty requires. But there are many close worlds. Will not the Safety Condition be too easily satisfied?

Take *Plane Crash*. There is a close world, through comparison with which, Passenger is made worse off by being denied admittance onto the plane. The Safety Condition is satisfied, so Passenger has a right against being denied admittance onto the plane. However, there is also another close world, through comparison with which, Passenger is made *better off* by being denied admittance onto the plane: through comparison of world 1 and world 2. Does this imply that Passenger has a right against Attendant that Attendant deny Passenger admittance onto the plane? It might appear so—the Safety Condition *has* been satisfied. This is one example of the Problem of Overgeneration.

Consider another case.

A&E. An unconscious Patient comes into A&E with a burst aneurysm. There is no chance that the aneurysm will stop bleeding spontaneously and so, without treatment, Patient will surely die. Surgeon can operate. The surgery is very serious, though all-things-considered beneficial.

¹⁷ Most difficult to swallow are going to be cases in which people work together to make harm very preempted. But perhaps there are things to say in reply to these cases, for example, that the victim has a right against the group acting.

While the case stipulates that there is no nomologically possible world in which Patient recovers without treatment (a possible world with our laws of nature), there might be a metaphysically close world in which the aneurysm stops bleeding. On the Lewisian view, all that would be required is a small miracle to stop the aneurysm from bleeding. Through comparison with this world, Patient would be worse off as a result of the operation: she will unnecessarily have gone through a very serious operation. Does this imply that Patient has a right against Surgeon performing the operation *because* there is a metaphysically close, though nomologically impossible, world in which the bleeding would stop?

I think the Problem of Overgeneration can be resisted. In what follows, I do so on several fronts. The general strategy is that we rule out worlds which would otherwise allow the Safety Condition to overgenerate rights by appealing to other features of closeness or to other necessary but insufficient conditions on rights.

Before that, note the Problem of Overgeneration does not arise if we prefer the Relevance Variant of the Safety Condition. Since there is no stipulation that close worlds are relevant, we need not worry that world 2 is very close in *Plane Crash*, since we can say world 2 is not relevant. Similarly, we need not worry about the metaphysically close world in which Patient's burst aneurysm stops bleeding. Obviously enough, one might still be worried about the view's epistemic circularity. So, let us see what can be said on behalf of the Safety Condition without moving to the Relevance Variant.

First, we may introduce what we can call the *Realism Condition*, according to which only those worlds that are nomologically possible count as close for rights.¹⁸ In *A&E*, the metaphysically possible world in which the aneurysm stops bleeding will not close, so the

¹⁸ See (McMahan 2002, 133–36).

Safety Condition will not be satisfied. However, the Realism Condition takes us only so far. It is of little use in *Plane Crash*.

Second, another feature of closeness that helps us deal with cases like *Plane Crash* is that the closeness relation is what we can call *nonreciprocal*—just because w_1 is close to w_2 when w_1 is our focus does not entail that w_2 is close to w_1 when w_2 is our focus.¹⁹ For example, though if the plane crashes there *is* a close world in which the plane does not crash, it is often *not* the case that, if the plane does not crash, there is a close world in which the plane does crash. Roughly, this is because a lot needs to go wrong for a plane *to* crash—meaning the world in which it does crash is far away. But it does not take a lot for a plane *not* to crash—meaning that the world in which the plane does not crash is close. While this does not explain why Passenger does not have a right that *she* be denied admittance onto the plane (since in that case, there is a close world in which the plane crashes), it does explain why, generally, *we* do not have rights to be denied admittance onto planes—usually, there is no close world in which the plane crashes. The Problem of Overgeneration is not as worrying as it might have seemed.

Third, recall the Safety Condition begins,

X has a right against Y that $Y \Phi$, only if (and *because*) Y 's not Φ -ing *causes* X to be worse off...

That Y needs to *cause* X to be worse off seems to get the cases right: Attendant causes Passenger to be worse off through being denied admittance onto the plane were the plane not to have crashed; Hitman₁ causes Victim to be worse off were Hitman₂ not present; and Shooter causes Target to be worse off were the risk to materialise. But, were Attendant to

¹⁹ For similar discussion, see (Lewis 1973, 50–52).

allow Passenger onto the plane, she does not *cause* Passenger to be worse off.²⁰ Given this, the Safety Condition is not satisfied, so Passenger will not have a right that Attendant deny her admittance onto the plane.

There may be problems with this causal restriction on the relationship between \mathcal{Y} 's Φ -ing and how X fares, as there might be cases in which \mathcal{Y} 's not Φ -ing does not cause X to be worse off, yet we think X has a right against \mathcal{Y} that she Φ . Without a theory of causation on the table, we are not able to fully assess this issue. Instead, let us move onto one final way to resist the Problem of Overgeneration: the Safety Condition is a necessary but insufficient condition on rights. And this means that X might not have a right that \mathcal{Y} Φ even if there is a close world in which \mathcal{Y} 's Φ -ing leaves X much better off.

So, what other necessary but insufficient conditions would need to be satisfied? Consider the following two popular moral asymmetries:

Doctrine of Double Effect. Other things being equal, it is harder to justify \mathcal{Y} intending to harm X than it is to justify \mathcal{Y} merely foreseeing that she will harm X .

Doctrine of Doing and Allowing. Other things being equal, it is harder to justify \mathcal{Y} doing harm to X than it is to justify \mathcal{Y} allowing harm to X .

Suppose we made DDE internal to rights. We would say that, all else equal, it is more likely to warrant a right if \mathcal{Y} intends the harm that would befall X were she to violate the putative duty than if that harm were merely foreseen. This means that a harm to X caused by \mathcal{Y} might be sufficient to warrant a right if that harm is intended by \mathcal{Y} , but insufficient if it was an unintended.

²⁰ This is not because omissions are not causal: (Lewis 1986, 189–93; Paul and Hall 2013, 173–214).

While both conditions are *ceteris paribus*, and stated generally, endorsing either (or both) helps resist the Problem of Overgeneration. Suppose Attendant were to admit Passenger onto the plane and the plane were to crash. Given the way that the case is stipulated, it is unlikely that Attendant *intends* the harm or *does* harm to Passenger. This helps explain why Passenger fails to have a right that Attendant deny her admittance onto the plane (notwithstanding my suggestion above that the Safety Condition is not satisfied in this variant in any case, because Attendant would not *cause* harm to Passenger were she to allow her onto the plane).

I have offered three dimensions along which Passenger might fail to have a right that Attendant deny her admittance onto the plane. First, Attendant would not cause Passenger to be worse off were she to allow Passenger onto the plane. Second and third, Attendant does not intend nor *do* the harm to Passenger. Taken with the two considerations above—that only nomologically worlds count as close for the Safety Condition and that the closeness relation is nonreciprocal—a picture is emerging of how we might respond to the Problem of Overgeneration.

5. Conclusion

I have introduced the Safety Condition as a reply to the Problem of Harmless Wronging for the Interest Theory. The problem is that of accommodating our intuitions that agents have rights against harmless wrongs, given a commitment to the Interest Theory. The Safety Condition's extensional accuracy and unified account of *why* agents are attributed rights across these different types of case of harmless wrongdoing is the primary virtue of the account that I have stressed in this chapter.

6. Acknowledgements

This chapter is based on my PhD thesis, so there are more people to thank by name than space permits—but thanks to everyone with whom I have discussed safety, and especially to those who have provided me with written comments on drafts of this paper, including two anonymous readers for this volume. Thanks to audiences in Cambridge, Edinburgh, Oxford, St Andrews, Boulder (and especially to Bodhi Melnitzer, my excellent commentator), Berlin, and to those who attended two online workshops at which this chapter was presented. Finally, thanks to my supervisors, Rowan Cruft and Theron Pummer, and to Stephanie Bowen, for their unfaltering support. This work was supported by an AHRC Scottish Graduate School for the Arts and Humanities doctoral training partnership studentship (grant no. 1804818) and a Society for Applied Philosophy Doctoral Scholarship.

7. Works Cited

- Bowen, Joseph. 2021. 'Review of Jonathan Quong, *The Morality of Defensive Force*'. *Ethics* 131 (3): 625–30.
- . Forthcoming. "'But You Could've Hurt Me!': Risk and Harm". *Law and Philosophy*.
- Bradley, Ben. 2009. *Well-Being and Death*. Oxford University Press.
- . 2012. 'Doing Away with Harm'. *Philosophy and Phenomenological Research* 85 (2): 390–412.
- Finkelstein, Claire. 2003. 'Is Risk a Harm?'. *University of Pennsylvania Law Review* 151 (3): 963–1001.
- Kamm, F. M. 2007. *Intricate Ethics*. Oxford: Oxford University Press.
- Lazar, Seth. 2017. 'Risky Killing: How Risks Worsen Violations of Objective Rights'. *Journal of Moral Philosophy*.
- Lewis, David K. 1973. *Counterfactuals*. Oxford: Basil Blackwell.
- . 1979. 'Counterfactual Dependence and Time's Arrow'. *Noûs* 13 (4): 455–76.
- . 1986. 'Causation'. In *Philosophical Papers. Vol. 2*, 159–72. Oxford: Oxford University Press.
- McMahan, Jeff. 2002. *The Ethics of Killing: Problems at the Margins of Life*. Oxford: Oxford University Press.
- Oberdiek, John. 2017. *Imposing Risk: A Normative Framework*. Oxford: Oxford University Press.
- Paul, L. A., and Ned Hall. 2013. *Causation: A User's Guide*. Oxford, New York: Oxford University Press.
- Quong, Jonathan. 2015. 'Rights Against Harm'. *Aristotelian Society Supplementary Volume* 89 (1): 249–66.
- . 2020. *The Morality of Defensive Force*. Oxford: Oxford University Press.
- Raz, Joseph. 1992. 'Rights and Individual Well-Being'. *Ratio Juris* 5 (2): 127–42.
- Smith, Martin. 2016. *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Tadros, Victor. 2016. *Wrongs and Crimes*. Oxford: Oxford University Press.
- Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Cambridge, Mass: Harvard University Press.
- Vossen, Bas van der. 2016. 'Uncertain Rights Against Defense'. *Social Philosophy and Policy* 32 (02): 129–45.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- . 2009. 'Replies to Critics'. In *Williamson on Knowledge*. Oxford: Oxford University Press.
- Zimmerman, Michael J. 2014. *Ignorance and Moral Obligation*. Oxford: Oxford University Press.