



Risky Killing

How Risks Worsen Violations of Objective Rights

Seth Lazar

Australian National University

sethlazar@gmail.com

Abstract

I argue that riskier killings of innocent people are, other things equal, objectively worse than less risky killings. I ground these views in considerations of disrespect and security. Killing someone more riskily shows greater disrespect for him by more grievously undervaluing his standing and interests, and more seriously undermines his security by exposing a disposition to harm him across all counterfactual scenarios in which the probability of killing an innocent person is that high or less. I argue that the salient probabilities are the agent's sincere, sane, subjective probabilities, and that this thesis is relevant whether your risk-taking pertains to the probability of killing a person or to the probability that the person you kill is not liable to be killed. I then defend the view's relevance to intentional killing; show how it differs from an account of blameworthiness; and explain its significance for all-things-considered justification and justification under uncertainty.

Keywords

risk – harm – rights – pro tanto wrongfulness

1 Introduction

Many believe that harming intentionally is objectively worse than doing so as a foreseen but unintended consequence of your action. In this paper, I argue for

Seth Lazar is Associate Professor and Head of School, at the School of Philosophy, Research School of Social Sciences, ANU. He works on the ethics of risk, self-defence, and war. His most recent book is *Sparing Civilians* (OUP, 2015). His current project, on deontological decision theory, is called *Duty under Doubt*.

a similar distinction between harms that are more and less foreseeable. Harming someone when it was more likely that your action would harm an innocent person is objectively worse than doing so when it was less likely.¹ Consider these two cases:

Demolition 1: Allie, a demolition worker, is about to destroy a building to clear a site for development. The building has been checked and there are no noises or movements. The probability someone is inside is very low – say 0.01. Knowing this, Allie demolishes the building. She kills Bruce, who was trapped inside through no fault of his own.

Demolition 2: Same as Demolition 1, except Allie can hear a voice calling from the building. It could be a TV, so the building might still be unoccupied; but that's unlikely. The probability of someone being inside is 0.8, say. Knowing this, she proceeds, and again kills Bruce.

There are two natural ways to assess these cases. The first suggests that Allie is more blameworthy in Demolition 2 than in Demolition 1. The second argues that (on some plausible assumptions) her action is wrong in light of her beliefs or evidence in Demolition 2, but not in 1. But both views agree that in light of all the facts, Allie's action is equally morally bad in either case.²

I think this is a mistake. Allie's action is objectively worse when the probability that she would wrongfully kill an innocent person was higher. Other things equal, riskier killings (and other harms) are more seriously wrongful than less risky killings (and other harms). Of course, many have argued that risks harm those exposed to them, and that we have rights against risks.³ Although I develop my own account of the wrong and harm in risk-imposition here, I will not criticize those alternative accounts – they may also be compatible with my central thesis. But almost nobody, to my knowledge, has noticed that these

1 For simplicity, I will focus throughout on killing. But everything I say applies to other harms as well.

2 See, for example, Sven Ove Hansson, 'In Defence of Deontic Diversity', *Journal of Logic and Computation*, (2015), pp. 1–19; Stephen Perry, 'Risk, Harm, Interests, and Rights', in Tim Lewens (ed.), *Risk: Philosophical Perspectives* (New York: Routledge, 2007), pp. 190–210; M.J. Zimmerman, 'Risk, Rights, and Restitution', *Philosophical Studies*, 128 (2006), pp. 285–311.

3 See, in particular, Claire Finkelstein, 'Is Risk a Harm?', *University of Pennsylvania Law Review*, 151 (2003), pp. 963–1001; John Oberdiek, 'Towards a Right against Risking', *Law and Philosophy*, 28 (2009), pp. 367–92; Stephen Perry, 'Risk, Harm, Interests, and Rights'; Judith Jarvis Thomson, *Rights, Restitution, and Risk: Essays in Moral Theory* (Cambridge, Mass.: Harvard University Press, 1986); M.J. Zimmerman, 'Risk, Rights, and Restitution'; Sven Ove Hansson, *The Ethics of Risk* (London: Palgrave Macmillan, 2013).

risks count against one's action objectively, even when it in fact results in a violation of rights.⁴

Section 2 gives a conceptual toolkit; Section 3 offers arguments for my view; Section 4 asks which kind of probability I have in mind. Section 5 explains how my arguments apply when we vary the probability of one's killing being *wrongful*, rather than the probability that one kills. Section 6 considers objections; Section 7 concludes.

2 Concepts

Rights are reasons, grounded in our moral status, protecting our interests. The right to life is the most important right, and we all start out enjoying its protection. But sometimes that protection can be weakened or lost, perhaps by waiver, or forfeiture, and certainly when it is necessary and proportionate to kill you to avert an unjustified threat, for which you are sufficiently responsible. In that case killing you does not violate your right to life: you are liable to be killed. When someone is not liable to be killed, following the literature in the ethics of harm, I will call her 'innocent'.⁵

Killing the innocent is pro tanto wrongful. This means that a weighty moral reason tells against that action: the victim, whose right is violated, has been wronged by the perpetrator. Pro tanto wrongfulness comes in degrees; there are more and less grave wrongs: contrast murder and manslaughter, for example. Sometimes, other reasons override our reasons not to kill the innocent. They can do so less readily when the pro tanto wrongdoing is graver. I focus

4 Well, not quite *nobody*. I apply the results of this argument in a companion paper on the ethics of war, and Mike Otsuka makes a similar observation, in passing, in Michael Otsuka, 'Risking Life and Limb: How to Discount Harms by Their Probability', in Nir Eyal, I Glenn Cohen, and Norman Daniels (eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective* (Oxford: Oxford University Press, 2015), pp. 77–93. See also Seth Lazar, 'Risky Killing and the Ethics of War', *Ethics*, 126 (2015), pp. 91–117.

5 See, for example, Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009). Although 'innocent' is a term of art, it is used advisedly. On my understanding, with one possible exception, the fact that one is innocent of responsibility for a wrongful threat is sufficient grounds for one not being liable to be killed in order to avert that threat. The exception: I suspect that some kinds of duress can render one's contribution to a wrongful threat morally innocent, i.e. blameless, but that one can nonetheless be liable to be killed in such cases. It is important to be clear, though, that innocence here refers exclusively to one's degree of responsibility for the wrongful threat that is to be averted. It does not invoke a more general property of one's character.

throughout on pro tanto wrongfulness, so will drop the ‘pro tanto,’ except where it is needed for emphasis. When I describe one action as ‘worse’ than another, I mean ‘more gravely pro tanto wrongful than.’ It’s also worth noting that one action might be thought worse than another in virtue of its consequences. My focus throughout is on non-instrumental pro tanto wrongfulness – the wrongfulness of ‘the act itself.’⁶

Objective probabilities – chances – are mind-independent features of the world, typically of the kind that the sciences aim to discover, like the half-life of Carbon-14, or the chance that a photon will head left or right in a double-slit experiment.⁷ Evidential probabilities measure the support a proposition receives from a body of evidence.⁸ Subjective probabilities are probabilities as assigned by a subject: they are credences in some proposition, or degrees of belief in it.⁹ Evidential and subjective probability both presuppose a limited epistemic perspective – either the agent’s evidence, or her beliefs. The primary probabilities in this paper must be epistemic in this sense. Bruce is in fact inside the building when Allie presses the plunger.

An act is pro tanto wrongful if a sufficiently weighty moral reason counts against it non-instrumentally (that is, not in virtue of its consequences). (Pro tanto) wrongfulness also comes in objective, evidential, and subjective varieties. When A ϕ s, ϕ ing is objectively wrongful if it is wrongful in light of all the facts that could bear on its deontic status.¹⁰ A’s ϕ ing is evidentially wrongful if it is wrongful in light of A’s evidence when she ϕ s. And it is subjectively wrongful if it is wrongful in light of what A believes when she ϕ s. Subjective and evidential standards can again be grouped under the epistemic heading.¹¹

6 Jonathan Bennett, *The Act Itself* (Oxford: Clarendon Press, 1995). Thanks to Victor Tadros for helping me see this point.

7 Rachael Briggs, ‘The Metaphysics of Chance’, *Philosophy Compass*, 5 (2010), pp. 938–52.

8 E.g. John Maynard Keynes, *A Treatise on Probability* (London: Macmillan, 1921); Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000).

9 E.g. F.P. Ramsey, ‘Truth and Probability’, in Antony Eagle (ed.), *Philosophy of Probability: Contemporary Readings* (Routledge, 2010), pp. 52–94; Bruno De Finetti, *Theory of Probability: A Critical Introductory Treatment* (Chichester: Wiley, 1990).

10 Notice, though, that ‘objective’ here does not mean (as it does with probabilities) mind-independence.

11 My aim in using this terminology is to emphasize the commonalities between the classification of probability and of wrongfulness. It would perhaps be more natural to speak of doxastic and evidential wrongfulness as subspecies of subjective wrongfulness, in contrast with the objective variety. But it would be too revisionary to group doxastic and evidential probabilities under the heading of subjective probabilities. Thanks to an editor for pressing me on this point.

We can now more precisely describe the dispute about Demolition 1 and 2. Objectivists about wrongfulness assert that, in light of all the facts, Allie's action is equally objectively wrongful in the two cases, since she in fact violates Bruce's right to life in both cases.¹² Subjectivists and evidentialists about wrongfulness argue that we should also (or perhaps only) assess Allie's action in light of her beliefs or her evidence.¹³ There are two prominent approaches. The first asks simply whether it was sufficiently likely that her action was wrongful.¹⁴ The second focuses on the expected value of her choices. Both will most likely agree that Allie's action is epistemically wrongful in Demolition 2, but not in Demolition 1.

I assess these cases differently: 'all the facts' relevant to the objective evaluation of Allie's action include facts about the epistemic probability (leave it open, for now, whether that probability is subjective or evidential) that her action would kill an innocent person. When that probability is higher, her action is *objectively* worse than when it is lower. In what follows, I defend the following principle:

Risky Killing (RK): Other things equal, when A's ϕ ing kills an innocent person, her action is objectively worse the higher the probability when she ϕ d that her action would kill an innocent person.

3 Respect and Security¹⁵

I have two arguments for RK. Both aim to show (a) that we have reasons not to impose risks on others, which (b) obtain even if those risks in fact result in harm. Other philosophers have focused, almost exclusively, on showing that we have reasons not to impose risks on others, when those risks *don't* result in harm.¹⁶ They say little, if anything, about what happens to those reasons when the risked harms come about. Let me emphasize this point. None of the authors cited above – nor others working in this area – have addressed at length

12 For this kind of view, see Judith Jarvis Thomson, *Rights, Restitution*.

13 See Holly M. Smith, 'Subjective Rightness', *Social Philosophy and Policy*, 27 (2010), pp. 64–110.

14 Jeff McMahan, *Killing in War*.

15 The arguments made in this section extend and complement those made in Seth Lazar, 'Risky Killing', which focuses exclusively on cases in which the probability of innocence varies, rather than the probability of harmfulness (see Section 5 below).

16 See footnote 3 above. The only exception is a brief discussion, in passing, in Michael Otsuka, 'Risking Life and Limb: How to Discount Harms by Their Probability'.

the question of the *objective* wrongfulness of risk-imposition, in cases in which the imposed risks result in harm. Most of the debate has focused on whether we have rights not to be subjected to ‘pure risks’ – risks that do not result in harms. This is an important question, but it is different from my question. I am interested in (a) only insofar as it is necessary to deliver (b). So I do not, in what follows, engage directly with the arguments of those other theorists: though they make some compelling points, their quarry is different from mine, so it makes sense to concentrate on putting forward my positive case, rather than knocking down the competition. What follows, then, is a novel account of how risks aggravate objective rights-violations. My first argument focuses on respect, my second on security.

Allie is a normal person, with normal capacities and beliefs. She knows people have moral standing, and rights to life, which create duties not to kill them. Her perception is unclouded and unmanipulated, and her habits of induction are reliable. She knows that acts have consequences, and so she acts, given her beliefs, to achieve outcomes that she desires. She is not compelled; her desires are her own. She knows that demolishing the building realizes an inconsequential commercial end, but will surely kill anyone inside. And if someone were inside, then only by extraordinary coincidence would they not be innocent.

Now consider Demolition 3, in which Allie is certain that Bruce is inside the building. Still, she blows up the building, killing Bruce. What can we infer from her action? She must either be wholly indifferent to Bruce’s right to life, or deny that he has one. Either amounts to a radical failure of respect for Bruce: his death gives her no more pause than those of the insects also consumed in the explosion.

Rights protect beings with moral status. Respect is part of this protection: to violate someone’s right is not merely to harm him, but to show disrespect for his moral standing. Some rights violations are more disrespectful than others. The more disrespectful, the graver the wrongdoing. When Allie knows that her action will violate Bruce’s right to life, for no significant end, her action is extremely disrespectful.

In Demolition 2, the probability that Bruce is inside the building is much higher than in Demolition 1. Allie’s action expresses disrespect in rough proportion to the probability that someone would be inside.¹⁷ When it is 0.8, as in Demolition 2, Allie’s action evinces much more disrespect than in Demolition 1,

17 At the very least, the relation between probability and disrespect is monotonically increasing. I say more on this below, and address some interesting complications raised by Yitzhak Benbaji and Victor Tadros, in Seth Lazar, ‘Response: Strengthening Moral Distinction’, *Law and Philosophy*, (Forthcoming).

when it is 0.01. Since it is more disrespectful, it is a graver wrongdoing, thus supporting RK.

One way to see this is to think of Allie's action from a decision-theoretic perspective. We can roughly infer the relative values that you attach to your decision's outcomes, if we know their probability of coming about. In Demolition 2, Allie is sure to realize the commercial benefits of destroying the building, but has a 0.8 chance of killing an innocent person. Call that commercial value B, and the value of not killing an innocent person V. By going ahead, Allie shows that she thinks B is no less valuable than 0.8V. Avoiding killing an innocent person is worth no more than 1.25 times the profits of destroying the building. This egregiously underestimates the importance of her victim's interests and standing.

In Demolition 1, by contrast, we can infer that $B \geq 0.01V$. So avoiding killing an innocent person is no more than 100 times more important than the profits from demolishing the building. Perhaps even this disrespects her victim, but it is clearly better than in Demolition 2.

One might object, here, that we cannot infer Allie's attitude to her victim from the probabilities on which she acted – her occurrent attitudes to her victim could obviously diverge from the attitude implied by her action. Suppose, for example, she regards him not with ill will, but with regret. However, my argument is not that Allie explicitly shows ill will to her victim. She *lacks* the proper respect. This need not be expressed in a positive attitude – *indifference* is sufficient.¹⁸ Even if Allie regrets Bruce's suffering, that she nonetheless goes ahead shows that she is unacceptably indifferent to his standing and status.

But what if Allie is solicitous of Bruce's wellbeing, thinking she is shepherding him to eternity in heaven? She still lacks the proper respect. Attitudes are not always transparent to the person who bears them.¹⁹ She might convince herself that she is acting from love for Bruce, but in fact her disrespect for him is untempered. More on this below.

The additional risk Allie takes in Demolition 2 aggravates her violation of Bruce's right to life. But it also contravenes another important interest of Bruce's: that in *being secure*.²⁰

18 For some useful reflection on indifference, see James B. Brady, 'Recklessness', *Law and Philosophy*, 15 (1996), pp. 183–200, 193.

19 Timothy Williamson, *Knowledge and Its Limits*, Chapter 4; James B. Brady, 'Recklessness', 187; Matt King, 'The Problem with Negligence', *Social Theory and Practice*, 35 (2009), pp. 577–95, 583.

20 Bruce plausibly has a right to security, which is contravened by Allie's action; however, since in this case the right would be grounded in her interest, it seems simpler to just focus on the additional harm involved in exposing Bruce to risk, rather than to make the further step to there being a right here.

Philip Pettit has recently introduced the idea of a *robustly demanding good*, to enjoy which one must enjoy a counterpart thin good not only in the actual world, but also in a range of counterfactual scenarios.²¹ For example, to be truly free others must not interfere with your choices not only in the actual world, but in counterfactual scenarios in which your preferences, or their dispositions towards you, are different. And a real friendship requires not only mutual favour between two people in the actual world, but continued mutual support across counterfactual scenarios. If one would ditch the other if he lost his money or his looks, then they are not truly friends.

I think there is a further robustly demanding good: *security*, the robust avoidance of *pro tanto* wrongful harm. To enjoy security, one must not only avoid wrongful harm in the actual world, but also do so across relevant counterfactual scenarios: those in which the victim does not get lucky. We are insecure to the extent that others make our avoidance of wrongful harm depend on luck.

Security and cognate goods are closely tied to contingent benefits such as peace of mind, or the ability to plan for the future.²² And in practice, since we cannot know the future, we cannot enjoy these goods unless we are secure. However, I think security is also non-instrumentally valuable, independently of these contingent benefits, for two reasons.²³

First, it is non-instrumentally worse to avoid wrongful harm merely through good luck, than to do so robustly, because the more you depend on luck, the less control you have over your life, and so the less autonomous you are. Autonomy is non-instrumentally valuable, so its constituent parts – such as control over whether your most important interests are satisfied – are non-instrumentally valuable too.

Second, one's security is often grounded in others' positive dispositions towards one. I am secure not merely through my own efforts, but because others are reliably disposed not to threaten me, to protect me against threats, to make me whole if this protection fails. I am secure because I am an object of concern

21 Philip Pettit, *The Robust Demands of the Good* (Oxford: Oxford University Press, 2015); Philip Pettit, *On the People's Terms: A Republican Theory and Model of Democracy* (Cambridge: Cambridge University Press, 2012). See also Nicholas Southwood, 'Democracy as a Modally Demanding Value', *Nous*, 49 (2015), pp. 504–521.

22 For a recent survey of such views, see Madeleine Hayenhjelm and Jonathan Wolff, 'The Moral Problem of Risk Impositions: A Survey', *European Journal of Philosophy*, 20 (2012), pp. 26–51, 31.

23 The next two paragraphs overlap with arguments made in Seth Lazar, 'Risky Killing'. For a different autonomy-based argument for a right against risk, see John Oberdiek, 'The Moral Significance of Risking', *Legal Theory*, 18 (2012), pp. 339–56.

for those around me. Even if I need never draw on others' concern, I am better off just by their having this positive disposition towards me. And I am worse off when they make their impact on my interests a mere matter of luck.²⁴ The higher the probability, from Allie's perspective, that her action will kill an innocent person, the more she makes Bruce dependent on luck for his survival.

Being secure brings us contingent benefits; it makes us freer, and it evinces our standing as equal members of the moral community, objects of others' concern, whose rights are not merely accidentally respected. But how does it apply to our cases? Bruce is in fact in the building. Does Allie therefore undermine his security to the same degree regardless of her probabilities? No. She makes Bruce's avoidance of wrongful harm more dependent on luck in Demolition 2, because she shows her readiness to proceed in all counterfactual scenarios for which the probability of someone being in the building is less than 0.8. In Demolition 1, by contrast, she shows herself ready to proceed only in scenarios where the probability is below 0.01. Even though Bruce is in fact in the building both times, he is more unlucky to be harmed in Demolition 1 than in Demolition 2, because there are fewer counterfactual scenarios in which the probability of killing a person is below 0.01 than there are in which it is below 0.8 (the latter is obviously a superset of the former). Conversely, to avoid harm in Demolition 2, he would have to be very lucky – the probability of his being in the building would have to be greater than 0.8.

Allie undermines Bruce's security to a greater degree in Demolition 2 than in Demolition 1. This aggravates her ultimate wrongdoing. Besides killing Bruce, she also renders him insecure. But might one object that Bruce's interest in security is undermined only when he is subjected to risks, but when those risks result in actual harm, a different interest is undermined? I think not. Even when the risk imposed results in harm (and so breaches a distinct interest), Allie still undermines Bruce's security, just as she would have done had he got lucky, and escaped harm. His being killed does not vitiate an objection against Allie that he would have had, had he survived. If Allie exposed Charlie to just the same risk as she did Bruce, at the same time, but Charlie miraculously survived, then Allie would have damaged both Bruce's and Charlie's security interests, but only Bruce's right to life.

Another objection: are there not other ways in which Allie's harming Bruce is more or less robust? Perhaps it is a fluke that he is there on that day, for instance. Of course, my focus here is on only the pro tanto contribution of the probability that she kills an innocent person. That other properties contribute as well does not undermine this claim. Less concessively, our goal here is to

24 See Seth Lazar, *Sparing Civilians* (Oxford: Oxford University Press, 2015).

evaluate Allie's action. If other factors contribute to or undermine Bruce's security, that might be bad for Bruce, but it is irrelevant to how we evaluate what *Allie* did.

Most would agree that *mere* risk imposition can be wrongful. Most would also agree that higher risks are, other things equal, worse. If driving through a housing estate at 60mph is wrong, doing so at 150mph is worse, even if nobody is hurt. If operating while tipsy is wrong, doing so while smashed is worse, even if you miraculously remove the tumor. The argument from security's central claim is that, even if your risky behavior *does* result in harm, these facts still determine the gravity of your wrongdoing.

A last thought. Why does RK focus on the probability that *an* innocent person would be killed? Shouldn't it be individuated to the victim? This is a tricky question, and I have not reached a settled view. I think that the degree of disrespect manifest in Allie's action depends on the probability that the building would contain an innocent person, not the probability that Bruce himself would be in there. I also think that Bruce's security is undermined on the same grounds. This is intuitively plausible: it doesn't seem to make any moral difference if we imagine that Allie knows that Bruce is one of 10 people who might be inside the building, each of whom is equally likely to be the one person who is inside. In that variation on Demolition 2 Allie's credence that her action will kill Bruce is only 0.08, but it still seems that she wrongs him just as much as she wrongs him when he is the only person who might be in the building. This intuitive judgment reflects, I think, the fact that moral theory should not give much weight to the specific identity of a wrongdoer's victims – what matters is not that Allie disrespected *Bruce*, specifically, but that she disrespected an equal member of the moral community, who turned out to be Bruce. One way to model this would be to posit that Bruce's complaint against Allie is made on behalf of all the people who could have ended up in Bruce's position.

4 Which Probabilities?

Bruce is inside the building, and we can suppose that the objective probability of killing him is the same in both cases. So the probabilities presupposed by the foregoing arguments must be either evidential or subjective. This means, in turn, they can vary depending on whose epistemic perspective they are assessed from, and indeed depending on when they are assessed. From Bruce's perspective, the probability that Allie's pressing the plunger will kill him is much higher than it is from Allie's perspective. Which one counts?

We cannot tell anything about Allie's attitude toward her victim from probabilities that are epistemically inaccessible to her. And only when she acts does she decisively reveal those attitudes. So the first argument must assess probabilities from Allie's perspective, when she acts.

The argument from security *could* use Bruce's probabilities, and indeed objective ones. This is especially true for security's contingent benefits. But Allie's probabilities also matter: her readiness to proceed when the probability of harming Bruce is 0.8 indicates a broader readiness to proceed should the probability have been at least as great. That is the respect in which his avoidance of wrongful harm is more dependent on luck in Demolition 2 than in Demolition 1.

To keep things simple then, and since they are relevant to both arguments, I will focus on Allie's probabilities here – although with the caveat that Bruce's probabilities might also be relevant to the argument from security, as well as the observation that there is interesting work to be done on exploring the moral implications of objective risks.²⁵

Next question: are the salient probabilities evidential or subjective? The disrespect argument focuses on the agent's attitude to the victim. If the probabilities are evidential, then there is no obvious connection between those probabilities and Allie's attitudes. Her credences might diverge from the evidential probabilities, or she might assign no probabilities at all. So the probabilities in question cannot be evidential.

One might counter that our actions can convey attitudes regardless of what we actually think, just as words can express meaning independent of the speaker's intention. Actions, like words, can be symbols in a referential system that communicates meaning in virtue of context and audience. This idea is familiar in etiquette: when an infant swears, the word might be thought rude regardless of her intentions. Perhaps one could argue, then, that evidential probabilities are what matters. But defending this view would require an account of the

25 Here it is important to note that the fact that Bruce was ultimately harmed does *not* necessarily entail that the objective chance that he would be harmed was 1, as Stephen Perry, for example, asserts (Stephen Perry, 'Risk, Harm, Interests, and Rights'). Nor need this claim depend on arguing that the world is indeterministic. Some theories of objective chance (in particular functionalist, emergent and counterfactual ones) are consistent with the truth of determinism. See Aidan Lyon, 'Deterministic Probability: Neither Chance nor Credence', *Synthese*, 182 (2011), pp. 413–32; Christian List and Marcus Pivato, 'Emergent Chance', *Philosophical Review*, 124 (2015), pp. 119–52; Toby Handfield, *A Philosophical Guide to Chance* (Cambridge: Cambridge University Press, 2012).

social construction of meaning, showing how actions fit within that structure. Such a task is beyond my expertise, but may be worth further exploration.

What about the argument from security? There again Allie's subjective probabilities are what matters. We cannot infer a disposition to repeat her action in other similar circumstances from evidential probabilities of which Allie is unaware. Perhaps she ought to know the evidential probability, and acts wrongly in virtue of failing to do so, and so on. But the argument from security says that she would kill Bruce in any counterfactual scenario where he is in the building, and the probability is no greater than 0.8 (in *Demolition 2*). Only if the probability is subjective can we infer from what she does in the actual world, what she would do in those counterfactual scenarios.

So, RK can work with subjective probabilities (though perhaps not only with them). But how should we understand those subjective probabilities? Should we simply focus on Allie's actual probabilities? Or should they be idealized in some way?

First problem: Allie's actual subjective probabilities might be self-deceivingly or self-consciously insincere. She might convince herself that she thinks the probability of killing an innocent person is 0.01, while she really knows it is higher; or she might pretend that she had the exculpatory belief to avoid censure. Ultimately, however, while identifying her sincere subjective probabilities might sometimes be difficult, clearly these are what matters for determining the severity of her wrongdoing.

Second problem: sometimes people's beliefs are sincere, but insane. Suppose Allie sincerely believes that human beings are invulnerable to fire and force, so that even if there were people inside the building, they would not be harmed. Her subjective probability that she will kill an innocent person is therefore 0. Insofar as she genuinely believes this, then Allie is insane, and the moral principles governing the behavior of, and response to, the insane are different from those that are appropriate to interactions between the sane. But if Allie is sane, then on reflection she will abandon such absurd beliefs.

RK should at least deploy Allie's sincere, sane subjective probabilities. But perhaps we should go further, and seek the subjective probability a reasonable person would have, given Allie's evidence. Suppose, for example, that she hears a sound coming from the building, which any normal person would think a human cry, but Allie sincerely believes to be a bushbaby's call, irrelevant to her subjective probability that blowing the building up will kill an innocent person. This is not crazy, because bushbabies do sound like humans when they scream. But it is very stupid – Allie is in central London (and not near the zoo). Which is salient to evaluating her action – her sincere, sanely held subjective probability, or what a reasonable person would think in her place?

A reasonable person's beliefs in Allie's place are relevant to that person's attitudes, but not to Allie's. Provided her mistake is sincere, sane, and not based on an implicit rejection of central normative tenets, the only plausible route to inferring her attitude from her action goes via her actual subjective probability. Sometimes her mistake will be so stupid that it must either be insincere, or insane. But when it is genuinely sane and sincere, we lack grounds to impute to Allie any attitude beyond what that implies.

What goes for the disrespect argument goes for the security argument too. We can infer what Allie would do in salient counterfactuals only from her sincere, sane subjective probabilities. Insincere probabilities are uninformative about her actions in other cases; if she is insane, then we cannot infer from her ϕ ing in one context that she would do so in other similar contexts. But if she made a sincere, sane mistake, such that her subjective probability is 0.01 when it should have been 0.8, then we can infer only that she would proceed in cases where her subjective probability does not exceed 0.01.

Suppose Allie sincerely, sanely believes the probability is 0.01, but has made a terrible mistake, and it is really 0.8; suppose Allie* sincerely, sanely believes the probability is 0.01, and has got it right. Does each equally undermine Bruce's security? In all worlds where Allie or Allie* correctly believes the probability to be ≤ 0.01 , she proceeds. But since Allie is prone to making mistakes, wouldn't she proceed in a further set of worlds where she incorrectly believes the probability to be ≤ 0.01 ? Perhaps, but Allie's being prone to make mistakes does not entail that she will kill Bruce across more counterfactual scenarios. There might be false negatives as well as false positives.

But perhaps what matters here are not Allie's subjective probabilities, but her dispositions, for which her subjective probabilities are just unreliable evidence. Consider Demolition 1*, in which Allie's sincere, sane subjective probability is 0.01, but she is disposed to demolish the building for any probability ≤ 0.8 . Does the endangerment of Bruce track her subjective probability, which indicates only her minimal readiness to proceed, or her disposition, which dictates her actual readiness to proceed in counterfactual scenarios?²⁶

This is a difficult question, and I am open to resolving it either way. Perhaps Allie's dispositions are more fundamental than her beliefs where security is concerned. Indeed, perhaps we should understand her beliefs simply as a particular kind of disposition. Nonetheless, I think we can helpfully draw a contrast here between the evaluation of the action and the evaluation of the agent. When evaluating the action in the way required by the disrespect and security arguments, we must confine ourselves to considering what Allie has

²⁶ Thanks to an associate editor for help here.

committed herself to, through her action. Her action in Demolition 1* indicates that she regards a 0.01 risk of killing an innocent person as an acceptable one to run for the sake of achieving her objectives. If she would in fact tolerate an even higher risk, that clearly speaks to the evaluation of her character. But it is not a property of the action itself, which we can take to be the physical act, and her reasons for doing it.²⁷

And while Bruce might legitimately be dismayed that Allie has this disposition, the grounds of his complaint against her (or the one we can hypothetically make on his behalf) should be confined only to what Allie has committed herself through her action. Consider this analogy: suppose that Alf assaults Rajiv. Alf assaulted Rajiv because he was drunk and looking for a fight. It also happens that Alf is a racist, and on another occasion he might well attack someone just like Rajiv out of a racist motivation. Nonetheless, on this occasion, that was not part of his motivation. While Rajiv can be dismayed by Alf's racism, those racist dispositions are not relevant to Rajiv's grounds for complaint against Alf in this case. I think something similar applies in Demolition 1*: Allie's reasons for acting took into account the 0.01 risk of killing an innocent person. When thinking about the damage to Bruce's security for which *this act* is responsible, we should appeal only to that revealed disposition, rather than Allie's underlying recklessness.

5 From Lethality to Liability

The probability that one's action kills an innocent person might depend not only on its likelihood of being lethal, but on the probability that one's victim is innocent. Consider:

Demolition 4: the building to be demolished is near a public square, where a large rally is taking place. Allie's radio suddenly reports that a sniper is firing at the rally, picking people off. The announcer says that the shots are coming from her building. She looks up and sees a man apparently aiming a rifle at the square. She checks with a telescope, confirming that impression. There are no police nearby, and while she delays

²⁷ What if Allie had a very precise disposition, and would only proceed if the probability was exactly 0.8, no more or less? This would be an odd stance for her to take, but clearly she could not claim that she would not proceed when the probability is 0.7 on grounds of concern for her possible victims, so her indifference to her victim is no different from if she had the more standard attitude. Thanks to an associate editor for raising this question.

people are dying. Her sincere, sane, subjective probability that this man – in fact Bruce – is killing innocent people is 0.99, so the probability of non-liability is 0.01. She blows up the building, killing him. To her horror, the radio continues to report sniper fire and deaths – Bruce was holding a camera with telephoto lens, which from Allie's position looked exactly like a sniper rifle.

Demolition 5: the same as *Demolition 4*, except the radio does not report where the gunshots are coming from and Allie lacks a telescope, so can only dimly see Bruce on top of the building. Her subjective probability that he is the sniper is 0.2, so her subjective probability that he is not liable to be killed is 0.8. She blows up the building, killing innocent Bruce.

In both cases Allie violates Bruce's right to life, but she shows him more grievous disrespect, and undermines his security more seriously, in *Demolition 5* than in *Demolition 4*: the contrast between the cases is similar to that between *Demolition 2* and *1*. Allie shows greater disrespect for her victim in *Demolition 5* than in *Demolition 4*, and she also renders him less secure, by showing her readiness to proceed in any counterfactual scenarios in which the probability that he is innocent is at least that high.

So, I think the basic idea of RK extends to cases involving variable risks that one's target is innocent. However, it is worth pausing to note a difference between this case and those we have been discussing so far. In my view, the wrongfulness of harming an innocent person, at least when that harm is intentional (and perhaps in general) is an increasing marginal function of the probability that person is innocent. That is: the likelier it is that the victim is innocent, the faster the wrongfulness of harming him increases. Where the risk in question is one of lethality rather than of innocence, my view is that the wrongfulness of the act is roughly proportionate to the probability that the act would be harmful.²⁸

Focusing on probability of innocence raises a further interesting question about the scope of RK. At present it focuses only on cases in which A's victim is in fact innocent. How does it apply if the victim is in fact liable to be killed?

Consider *Demolition 5**, which is identical to *Demolition 5* except that, in fact, Bruce was the sniper, and Allie successfully though riskily stops his killing spree. Contrast this with *Demolition 4**, the same as *Demolition 4*, except that

28 I argue for this point at length in Seth Lazar, 'Response: Strengthening Moral Distinction'; Seth Lazar, 'In Dubious Battle: Uncertainty and the Ethics of Killing', *Unpublished ms.*, (2016).

Bruce is the sniper, and is liable to be killed. Does Allie wrong Bruce more seriously in the case where she takes the greater risk, and gets lucky?

The argument from disrespect supports the constrained formulation of RK. It shows that Allie's violation of Bruce's right to life can be aggravated when accompanied by an implicit indifference to her victim's standing and status. If Bruce is in fact liable, then his right to life is not violated, so Allie's disrespect has nothing to aggravate.

The argument from security, however, might support the opposite conclusion. Perhaps even though Bruce lacks the protection of his right to life, it remains either disproportionate or unnecessary for Allie to kill him in a way that also undermines his security.

The first possibility – that killing Bruce and undermining his security is disproportionate (though merely killing him would be proportionate) – does not seem right. The permission to kill someone to avert a threat that they pose is typically quite expansive. Suppose that in Demolition 5* Bruce survives the initial blast but is buried in the rubble, having suffered massive trauma. He survives for several hours, in agony, before finally dying. This additional suffering does not render killing him to save his victims impermissible; nor would knowing that he would suffer this way render it impermissible *ex ante*. To avert the threats that he poses, it is proportionate to kill him painlessly, but also proportionate to kill him in a way that involves excruciating pain if no less harmful alternative is available. I think that the further harm suffered by undermining Bruce's security also fits within the proportionality umbrella, justified by averting the threat Allie is trying to avert (although this is contingent on the magnitude of that threat).

However, even proportionate harms are impermissible if they are unnecessary – that is, if Allie could avert the threat without causing Bruce excruciating pain, then it is impermissible to do so in the more painful way.²⁹ Similarly, if she could avert the threat without undermining Bruce's security, then he retains his right to have his security maintained. If she has taken undue risks when she could have killed Bruce without undermining his security, her risky conduct is wrongful.

Ultimately I am unsure whether those who are liable to be killed can justifiably complain about their security being undermined. I therefore proceed with RK in its current form, with the caveat that it could perhaps be expanded to include liable as well as innocent victims.

29 See Seth Lazar, 'Necessity in Self-Defense and War', *Philosophy & Public Affairs*, 40 (2012), pp. 3–44.

6 Objections: Intention, Blame, Justification

It is worth addressing four further objections to my argument in detail. Each focuses not on the details of my case for saying that risks matter, but specifically on the thesis that risks aggravate otherwise wrongful acts. The first presents a putative counterexample, while the rest are deflationary, arguing that we can explain the same intuitions that RK covers by other, more familiar means.

Before Section 5, all my arguments focused on unintended killing. One might object that when killing is intentional, the distinctions drawn by RK are no longer relevant. Suppose Allie aims her six-shooter at Bruce, spins the barrel and pulls the trigger. The barrel stops at a full chamber, and Bruce is killed. Does the number of bullets really affect the gravity of her wrongdoing?³⁰

I think it does. Allie wrongs Bruce more gravely when she pulls the trigger more riskily. I think this would be clear if she got lucky, and the barrel stopped at an empty chamber – when Bruce survives, Allie has surely wronged him more gravely by putting in three bullets than by putting in one. Why should her successfully murdering Bruce diminish this complaint? Suppose the barrel has 1,000 chambers. Allie spins it again, it lands on a full chamber and she kills Bruce. Surely she has wronged him more gravely if she put 999 bullets in the barrel, than if she put in only one? I see no reason to deny the application of the arguments from disrespect and security to such cases.

Moreover, I have an error theory for the critic's intuition: a malevolent intention can drown out an action's other wrong-making features. Intending to {kill an innocent person for no reason} is paradigmatically wrong; acting on that intention is never justified. It is invidious enough that it can drown out or swamp other morally relevant considerations, so that even though they do in fact count, their impact is hard to detect. That is why some might feel that RK has little pull in the Russian roulette cases. By contrast, the intention to [kill someone who might be liable, to achieve some great good], is still very serious, but is clearly not as egregiously objectionable as the absolutely proscribed intention, hence will not drown out other morally relevant considerations to the same degree.

The next objection claims that I have misidentified my target. Risk-taking is relevant only to the evaluation of persons, not of actions. A riskier killing is more blameworthy than one that is less risky, but the killing itself is equally wrongful.

30 Derek Parfit, for one, thinks not. Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), 156ff.

First, even if this is a plausible objection to the respect argument, it does not undermine the argument from security, which says that Allie's risky behavior inflicts additional harm on Bruce. This must matter for the evaluation of her action, not only of her *qua* agent.

I also think this fails as an objection to the disrespect argument. The central claim is that the data that argument appeals to are fully explained by blameworthiness, so that any application to the evaluation of action is otiose. This invites three responses.

First, it is unsurprising that evaluations of blameworthiness and pro tanto wrongdoing travel together: how much we blame an agent in part depends on the gravity of her wrongdoing; we blame a murderer more than a mugger, for example, even if they are equally responsible for their crimes. Unsurprisingly, then, when Allie's wrongdoing is graver, because it is more disrespectful, she will often also be more blameworthy.

Second, though these evaluations often travel together, they sometimes come apart. Disrespect can aggravate pro tanto wrongdoing, without warranting increased blame; sometimes agents can be more blameworthy without aggravated wrongdoing.

Suppose Allie has been drugged without her knowledge and against her will; the drug makes her much more reckless than she would otherwise be; she then proceeds as in Demolition 2, blowing up the building despite a 0.8 subjective probability that there is an innocent person inside. Her action is as gravely wrong as in the original version of Demolition 2; it expresses the same degree of disrespect. And yet she is probably blameless, both for the action itself and for the disrespect she has shown the victim.

Or suppose Allie is disposed to press the plunger even if the probability that the building is occupied is 0.8, though in fact it is 0.01. Allie is blameworthy for her bad dispositions, even when they do not match up with pro tanto wrongful actions. The disrespect relevant to the wrongfulness of her action, however, is just that conveyed by her acting on the 0.01 subjective probability. Her actual readiness to proceed even if the probability is 0.8 is a property of Allie, but not a property of her action; her minimal readiness to proceed if the probability is up to 0.01 is a property of the action (and a property of Allie). Or so I argued above.

Lastly, suppose Allie's probability that someone is inside is 0.8, but the building is empty. There are no relevant counterfactual scenarios in which anyone would have been inside. This means nobody can say that Allie disrespected them, because nobody had any prospect of being in the building (equally, nobody's security is undermined). Again, Allie is blameworthy, but has not wronged anybody, so judgments of blameworthiness and pro tanto wrongfulness come apart.

Third response: the critic says that the data that explain blameworthiness and disrespect are irrelevant to permissibility; but surely if Allie knows that ϕ ing disrespects Bruce more than ψ ing, that is a reason to ψ rather than to ϕ .³¹ But if this is right, then we need some account of how blameworthiness is relevant to permissibility, not only to evaluating agents – just in case an action is more blameworthy, we have more reason to avoid it, for example. But if we offer that account, then the critic's worry is merely terminological. Where I speak of pro tanto wrongdoing, the critic speaks of blameworthiness, but insofar as they are both relevant to the evaluation and guidance of action, we are referring to the same things.

The third objection questions whether the arguments from disrespect and security work for risky conduct that is all things considered justified. Sometimes it is permissible to kill innocent people in order to save, or in the course of saving other people's lives. If infringing another's right to life is all things considered justified, does risk-taking aggravate that right infringement? Consider Demolition 1# and 2#, the same as Demolition 1 and 2 except that if Allie does not demolish the building immediately, it will collapse onto a school, killing scores of children. Make the numbers such that, even if she knew Bruce was inside, and not liable, she could permissibly proceed. Hasn't she given his interests and standing their due consideration? And how can he complain about the incursion on his security, when the good thereby achieved is so great?

How we answer this objection depends on how we understand the all things considered permissibility of imposing risks. One understanding can be quickly set aside. Sometimes the imposition of risks is justified as part of a fair distribution within society. Ordinary human life is impossible without imposing risks on others, so to avoid paralysis we are permitted to impose some risks on one another, provided we perform our duties of care.³² The paradigm case is driving: we can permissibly impose risks on others provided we are properly licensed, drive only when fit to, maintain our vehicles, and so on. When we impose risks on others that are justified in *this* way, they do not convey disrespect, or problematically undermine people's security. We would all consent to being exposed to such risks if given the opportunity. However, this is not the sort of all things considered justification that is relevant here: the appropriate social distribution of risks does not plausibly ground a right to run significant risks of killing others, only to (more or less reciprocally) impose some small risks on them.

31 I assume that both ϕ ing and ψ ing would be worse than just doing one or the other.

32 See e.g. Sven Ove Hansson, 'Ethical Criteria of Risk Acceptance', *Erkenntnis*, 59 (2003), pp. 291–309.

More interesting is lesser evil justification in extremis: cases in which one can permissibly infringe another's right, to bring about some greater good. In such cases I think that even if it is all things considered permissible to kill an innocent person, doing so more riskily aggravates the pro tanto wrongdoing he suffers. Allie still disrespects Bruce, as an individual, and she still problematically undermines his security, even if doing so is all things considered justified.

To start with, if it is all things considered permissible to *knowingly* infringe someone's right to life, then of course it will be permissible to do so more or less riskily. If ϕ ing is all things considered permissible, and ψ ing is pro tanto better than ϕ ing, then of course ψ ing will be all things considered permissible in the same scenario. This tells us nothing about whether and to what extent ϕ ing is worse than ψ ing.

I have claimed that when A riskily ϕ s, infringing B's right to life to bring about some outcome O, at least three classes of reasons tell against the action: the mere fact, R, that ϕ ing violates B's right to life; the disrespect, D, that aggravates that wrongdoing; and the degree to which B's security is undermined, S. If ϕ ing is riskier than ψ ing, then my view says that the D-reasons and S-reasons should tell more strongly against ϕ ing than against ψ ing. My critic responds that when the O-reasons (the reasons to bring about the outcome) override the R-reasons, then the D- and S-reasons are not simply overridden, they are silenced, or undermined – they no longer tell against ϕ ing. And if the O-reasons are enough to override the R-reasons, then it matters not whether A ϕ s or ψ s, because the S- and D-reasons are silenced.

I think this is a mistake. The S- and D-reasons are not undermined, but overridden; they retain their force. But even if I'm wrong about that, I think that the riskier A's ϕ ing is, the weightier the D- and S-reasons, *and the harder it is for the O-reasons to undermine them in this way*. Where ϕ ing is riskier than ψ ing, the O-reasons must be weightier in order for ϕ ing to be permissible than for ψ ing to be permissible. So, even if the D- and S-reasons are undermined, they push up the weight that the O-reasons must have for ϕ ing or ψ ing to be all things considered permissible.

To see this for the D-reasons, suppose that unintentionally killing one person as a side-effect of saving five is permissible. If you kill one as a side-effect of saving one, part of the reason that this is impermissible is the element of disrespect that you show your victim. If you kill one, saving two, perhaps the disrespect diminishes, but it still contributes to making your action impermissible. If you kill one and save five, perhaps the disrespect is undermined, but it still helped force up the outcome reasons, to make them weighty enough to render killing one all things considered permissible.

The last objection is similar in spirit to the objection from blame: it tries to more parsimoniously explain the data that grounds RK, by appealing to an account of epistemic permissibility (recall that epistemic permissibility includes evidential and subjective permissibility). For example, the critic might say: Allie acts subjectively permissibly in Demolition 1, and subjectively impermissibly in Demolition 2. That's all there is to it! There is no need for any odd claims about how her beliefs affect the objective evaluation of her action.

To respond, I will show how judgements of *pro tanto* wrongfulness can affect both objective and epistemic permissibility, in a manner that makes sense only if RK is true. I consider objective permissibility first, then the two variants of epistemic permissibility. I take no sides as to which sense of permissibility, if any, is more important than the others.

RK is relevant to objective permissibility in three distinct ways. First, suppose that A can choose between killing B and killing C, to achieve the same objective O. Suppose that achieving O is important enough to justify knowingly killing an innocent person. Neither B nor C is in fact liable. However, A's subjective probability that B is innocent is higher than his subjective probability that C is innocent. Even though achieving O could justify killing an innocent person, A is permitted only to kill C. Killing B would achieve her objective through action that was riskier than killing C, hence according to RK is unnecessarily wrongful.

Second, suppose O is a bad objective, which could not justify any killing whatsoever. I think A still acts *more* wrongfully if she kills B than if she kills C to achieve O. This is the realm of second-best oughts. Given that she is going to do something impermissible, she ought to do the least wrongful of the available actions, which is the less risky one. Of course, she really ought not to kill people to achieve that bad objective, but if she will nonetheless proceed, she should kill C rather than B.

Third, in virtue of RK, and its implications for the *pro tanto* wrongfulness of killing C and B, A's goal might be objectively valuable enough to justify killing C, but not B. If the probability that B is liable is lower than the probability that C is liable, then killing C is less risky than killing B. If it is less risky, then it is easier to justify. So O might be valuable enough to justify killing C, but not B. This is the most important role for RK. I think it potentially has serious implications for the morality of self-defense and war, which I gesture at in the conclusion, but discuss in detail elsewhere.

Next, epistemic permissibility. The first salient sense is as follows: you are epistemically permitted to do whatever, from your perspective, is sufficiently likely to be objectively permissible. An alternative view would use 'most likely' instead.

The manifold problems with this understanding of epistemic permissibility, and cognate understandings of ‘ought,’ have been discussed at length by others.³³ Here I simply consider the implications of RK for those, if there are any, who still construe epistemic permissibility this way.

We can start with cases where killing B or C will achieve O, which is sufficiently valuable to make knowingly killing an innocent person objectively permissible. As before, B is less likely to be liable than C. To work out what is epistemically permissible, A must determine what is objectively permissible in each outcome, and then assign probabilities to those outcomes. Table 1 lists the salient outcomes.

When either one or both are liable (scenarios 1, 3, and 4), matters are simple: killing either is objectively permissible if he is liable and objectively impermissible if the other is liable, and he is not. And what of scenario 2, in which neither is liable? Here RK makes an important difference. Recall that O is valuable enough to objectively justify knowingly killing an innocent person. Without RK, we could therefore conceivably think that killing either is objectively permissible, or perhaps even that killing B is objectively permissible, killing C objectively impermissible. But we can infer from RK and the details of the case that it is objectively impermissible to kill someone more riskily to achieve a given objective, no matter how valuable that objective, when you could have achieved the same result through a less risky killing. So RK tells A she is objectively permitted only to kill C in scenario 2.

As a result, killing B is epistemically permissible only if it is sufficiently likely that scenario 1 or 4 obtains; killing C is epistemically permissible only if it is sufficiently likely that scenario 1, 2, or 3 obtains. If RK were false, by contrast, killing B would be epistemically permissible if it were sufficiently likely that 1, 2, or 4 obtained. So, the truth of RK makes it less likely that killing B will be

TABLE 1

	B Liable?	C Liable?	Kill B objectively permissible?	Kill C objectively permissible?
1	Yes	Yes	Yes	Yes
2	No	No	No	Yes
3	No	Yes	No	Yes
4	Yes	No	Yes	No

33 E.g. Frank Jackson and M. Smith, ‘Absolutist Moral Theories and Uncertainty’, *Journal of Philosophy*, 103 (2006), pp. 267–83.

epistemically permissible (whether it will depend on the probabilities, and where the threshold for 'sufficiently likely' falls; or alternatively on whether given a number of options that are all sufficiently likely, one ought to choose the option most likely to be objectively permissible).

What of cases in which O is not sufficiently valuable to make killing an innocent person objectively permissible, whatever the probability that he is not liable? In these cases, RK still gives us an additional reason not to kill B, since we know that killing B if he is innocent is more seriously wrongful than killing C if he is innocent. This too will factor into our judgement of epistemic permissibility.

The same is true when O is sufficiently valuable to make killing an innocent person objectively permissible provided there is some probability – say, 0.5 – that he is liable to be killed. Again, RK can make the difference to epistemic permissibility by virtue of ruling out the possibility, in outcome 2, that killing B rather than C is objectively permissible. It can also intensify A's reasons to kill C rather than B – not only does she thereby run the smallest risk of killing an innocent person, but also, even if neither turns out not to be liable, she acts objectively permissibly. I return to this last point below.

The remaining sense of epistemic permissibility is, I think, the most interesting for RK. On this view, an act is epistemically permissible if it maximizes expected value.³⁴ In other words, we assign probabilities given that one ϕ s to the states of the world on which the various possible outcomes of ϕ ing depend, then assign values to those outcomes, multiply the probabilities by the values, and choose the act for which the sum of those products is greatest.

If RK is true, then killing an innocent person on a 0.8 probability that he was liable is objectively better than killing an innocent person on a 0.2 probability that he was liable. We must therefore build this into our expected value calculations. We can use numbers to illustrate the point – remembering that this is just an artificial heuristic.

Suppose the probability that C is liable is 0.5 and that B is liable is 0.1. As I noted above, I think it highly likely that the function from probability of innocence to wrongfulness increases at the margin: killing someone when there was a 0.9 probability that he was innocent is worse than killing him when the probability was 0.8 by a greater degree than would be the case if the probabilities had been 0.2 and 0.1 respectively. For present purposes, however, to keep things simple I will assume that the function is linear and increasing. So, suppose that the RK-weighting works like this: intentionally killing an innocent

34 As I argue in Seth Lazar, 'Deontological Decision Theory and Agent-Centred Options', *Ethics*, (Forthcoming); Seth Lazar, 'In Dubious Battle: Uncertainty and the Ethics of Killing', this is at most as a sufficient condition for epistemic permissibility, not a necessary one.

person when you are certain they are innocent has -100 disvalue; intentionally killing an innocent person when you are certain they are liable is -50 ; between those extremes the gravity of wrongdoing varies inversely proportionately with the probability of liability. So the disvalue of intentionally killing an innocent person who is 0.5 likely to be liable is -75 , and the disvalue of intentionally killing an innocent person who is 0.1 likely to be liable is -95 . The expected disvalue of killing C, *pro tanto*, is $0.5(-75)+0.5(0)$ (killing C if he is liable is not disvaluable). So: -37.5 . The expected objective disvalue of killing B, *pro tanto*, is $0.9(-95)+0.1(0) = -85.5$. Clearly, then, if A is permitted to kill either, she is permitted only to kill C.

One might object that RK makes no difference – if we ignore the RK-discount, the expected disvalue of killing C is $0.5(-100)+0.5(0) = -50$ and of killing B $0.9(-100)+0.1(0) = -90$. So the two calculations produce the same ranking. However, the cardinal difference between them matters! One could kill C to achieve objectives that would be insufficiently valuable without the discount. If killing C is sure to achieve O, with a resulting expected value of -40 , then it is permissible if RK is true, impermissible if not. This is all the more important if I'm right that the function from probability of innocence to wrongfulness is non-linearly increasing, since the more likely it is that one's target is innocent, the more steeply the wrongfulness of one's action increases.

One last thought: this approach to epistemic permissibility is risk-neutral. It gives additional weight neither to the best possible outcomes, nor to the worst ones. This is the standard approach to rational decision-making under risk, and is surely defensible at a first pass. However, one might think that we should build some significant risk-aversion into our theorizing.³⁵ In particular, nonconsequentialists who prioritize rights might do so by giving extra weight to the scenario where neither B nor C is liable. For this sort of risk-averse reasoning, which determines what is epistemically permissible for A to do by giving additional weight to the worst-case scenario, RK would be especially useful, since it tells us how to distinguish between more and less objectively *pro tanto* wrongful killings of innocent people.

7 Conclusion

I have argued that riskier killings of innocent people are, other things equal, worse than less risky killings. I grounded these views in considerations of

35 John Maynard Keynes, *A Treatise on Probability*, 360; see also Lara Buchak, *Risk and Rationality* (Oxford: Oxford University Press, 2013).

disrespect and security. Killing someone more riskily shows greater disrespect for him by more grievously undervaluing his standing and interests, and more seriously undermines his security by instantiating in your action a disposition to harm him across all relevantly similar counterfactual scenarios in which the probability of killing an innocent person is that high or less. I argued that the salient probabilities had to be the agent's sincere, sane, subjective probabilities, and that RK is equally relevant – though should be applied somewhat differently – whether your risk-taking pertains to the probability of killing a person or to the probability that the person you kill is not liable to be killed. I have defended the view's relevance to intentional killing; showed how it differs from an account of blameworthiness; explained its significance for all things considered justification; and schematically set out how it contributes to judgments of objective and epistemic permissibility. In a companion paper to this one, I address in detail the implications of RK for permissible killing in self- and other-defense, and especially in war.³⁶ Here it must suffice to make one comment.

Nonconsequentialist theorists of permissible harm typically affirm that the intentional violation of another's right to life is, if not absolutely impermissible, permissible only to avoid an unusual and severe catastrophe. Intentionally killing the innocent can be permissible in theory, but cases in which it is permissible are so rare as to be practically irrelevant. This generates serious problems for the plausibility of those views when applied to the ethics of self-defense and of war, in circumstances of uncertainty (which are of course endemic). In self- and other-defense, we often cannot know whether the person whom we have to kill to avert some threat is truly liable to be killed – we cannot know for sure the facts about causal and moral responsibility on which such judgments rest. As Frank Jackson and Michael Smith have argued, this poses serious problems for nonconsequentialists: if intentionally killing an innocent person is so terribly wrongful, how could I permissibly run even a small risk of committing such an egregious crime to achieve something as relatively inconsequential as saving a single life?³⁷

As I have argued elsewhere, this problem is compounded and multiplied many thousands of times over in the ethics of war.³⁸ On no plausible theory of

36 Seth Lazar, 'Risky Killing'.

37 Frank Jackson and Michael Smith, 'The Implementation Problem for Deontology', in Barry Maguire and Errol Lord (eds.), *Weighing Reasons* (Oxford: Oxford University Press, 2015), pp. 279–91.

38 E.g. Seth Lazar, 'The Responsibility Dilemma for *Killing in War*: A Review Essay', *Philosophy & Public Affairs*, 38 (2010), pp. 180–213.

liability to be killed will all the combatants whom we need to kill to win an otherwise justified war be liable to be killed. While we cannot usually tell which are liable and which are not, we know for sure that some significant proportion of the people we intentionally kill will *not* be liable to be killed. What could possibly justify running on such a massive scale the risk of committing wrongs that, even individually, could be permissible only to avert a rare and severe catastrophe? The difficulty of knowing whether one's targets are liable, combined with the near-absolutism of mainstream nonconsequentialist ethics, seems to drive inexorably towards a form of personal and collective pacifism.

RR pushes in the other direction: it says that some intentional killings of innocent people are objectively worse than others, and in particular that they are worse when they involve greater risks. If there is a high probability that your target is liable to be killed, then, *objectively*, killing him is less wrongful than if the probability of his being liable is lower. This may well mean that we can permissibly use lethal force against people of whose liability we are uncertain, to achieve the more prosaic ends for which self-defense and war are usually thought justified.³⁹

39 This paper was written with the support of the Australian Research Council, award DE130100811. I presented ancestors at the ANU and at the University of Adelaide. My thanks to all those who helped me improve the arguments of the paper, and in particular to the associate editor for this journal, for some particularly helpful comments.