

Proportionality & Uncertainty

Abstract

A military commander knows it is wrong to use *disproportionate* force in war. But she is uncertain whether a proposed bombing mission would be disproportionate. How should she proceed? More generally: how should considerations of proportionality enter into our decision-making about the use of force under conditions of uncertainty? Patrick Tomlin has recently developed what is, to date, the most careful and sophisticated answer to this question. In this note I argue that Tomlin's account, however, fails to do justice to the moral considerations underlying our concern that harm be proportionate, before proposing what I take to be a better answer.

1 Introduction

There is a strong presumption against harming others without their consent. Non-consensual harm may be imposed only under special circumstances — for example, as punishment for criminal wrongdoing, or to stop an aggressor from harming someone. But even in these special circumstances the use of harm is subject to various constraints. One of the most important and well-recognized of these constraints is the so-called

proportionality constraint. The fact that punitive or defensive harm would be "disproportionate" often makes it wrong to impose that harm.

Philosophers have had a great deal to say about this constraint in recent years.¹ This work, however, has predominantly focused on cases in which agents operate with full knowledge of their options and the consequences of those options. Human agents, of course, never operate with such knowledge. We must make decisions whether to harm in light of the limited information we have. And as such it is important that we have a theory that tells us how to take considerations of proportionality into account in light of that information, and not simply in light of the information we would have were we omniscient.

Patrick Tomlin has recently contributed to this lacuna in the literature.² He has provided what is, to date, the most careful and sophisticated attempt to articulate how we should make proportionality decisions under uncertainty. The following note is a critical response to this attempt. I argue that, while Tomlin identifies a serious problem with certain accounts, he fails to properly diagnose this problem — which in turn spells trouble for his alternative account. I propose a better way of accommodating the moral considerations underlying the importance of proportionality in our

¹See, for example, David Clark, "The Demands of Necessity," *Ethics* 133 (2023): 484-487; Kai Draper, "Necessity and Proportionality in Defense," in *The Ethics of Self-Defense*, eds. Christian Coons and Michael Weber (Oxford: Oxford University Press, 2016): 171-184; Thomas Hurka, "Proportionality in the Morality of War," *Philosophy & Public Affairs* 33 (2004): 34-66; F. M. Kamm, *Ethics for Enemies: Terror, Torture, and War* (Oxford: Oxford University Press, 2011): 130-156; Jeff McMahan, "Self-Defense and Culpability," *Law and Philosophy*, 24 (2005): 766; McMahan, "Proportionality and Time," *Ethics* 125 (2015): 1-25; McMahan, "The Limits of Self-Defense," in *The Ethics of Self-Defense*, eds. Christian Coons and Michael Weber (Oxford: Oxford University Press, 2016): 206-210; Jonathan Quong, *The Morality of Defensive Force* (Oxford: Oxford University Press, 2020): 96-123; Uwe Steinhoff, "Proportionality in Self-Defense," *The Journal of Ethics* 21 (2017): 263-289; Victor Tadros, *The Ends of Harm* (Oxford: Oxford University Press, 2011): 175-181, 331-360; Tomlin, "Proportionality in War: Revising Revisionism," *Ethics* 131 (2020): 34-61; Suzanne Uniacke, "Proportionality and Self-Defense," *Law and Philosophy* 30 (2011): 253-272.

²Patrick Tomlin, "Subjective Proportionality," *Ethics* 129 (2018): 254-283.

decision-making under uncertainty.

Overview. §2 summarizes Tomlin's account. §3 unpacks the shortcomings of Tomlin's account and points the way towards a better account. §4 concludes.

2 Tomlin on Subjective Proportionality

There is an intuitive distinction between what we might call "objective" and "subjective" proportionality. Consider:

Illusion. A soldier sees what appears to be an enemy combatant on a distant hill. She takes extensive precautions to confirm that the person is armed and that he poses an imminent threat before she opens fire, killing him. As it turns out, however, the intense desert heat created an optical illusion: her target was an unarmed and innocent shepherd.

Lucky Kill. A soldier sees an enemy combatant on a distant hill, taking shelter near what appear to be three innocent bystanders. She knows that the enemy combatant will kill her if she doesn't kill him first, but she can eliminate the threat only by throwing a grenade that she knows will also kill the three bystanders. She throws the grenade. Unbeknownst to her, however, the three bystanders are not actually bystanders, but enemy combatants who themselves posed an imminent and lethal threat to her.

In *Illusion*, the soldier uses force that is objectively disproportionate — that is excessive in relation to the threat actually posed by the shepherd. But there is clear sense in which this force is proportionate *in light of the soldier's evidence*. Her use of force is objectively disproportionate but subjectively proportionate. Likewise, this distinction seems apt in the case of

Lucky Kill. The soldier uses force that is subjectively disproportionate but objectively proportionate. It is excessive in light of her limited evidence, but not in light of the facts.

Much has been written on the nature of objective proportionality; much less on subjective proportionality. This is unfortunate. As Patrick Tomlin notes, an account of subjective proportionality is important for at least two reasons. First, "since it concerns how agents should perform proportionality calculations from within the epistemic position they in fact inhabit, such an account is the starting point for developing a proportionality standard that can be action-guiding."³ Second, it is subjective proportionality, and not objective proportionality, that is most relevant to the fittingness of praise and blame and to assignments of culpability.⁴

So how should we think about subjective proportionality? More precisely: how should considerations of proportionality enter into our decision-making about how to use force under conditions of uncertainty?

2.1 The Simple View

Tomlin takes as his foil what might seem the most "obvious way to move from objective proportionality to subjective proportionality." Call it the

SIMPLE VIEW. A use of force is subjectively proportionate if and only if the expected harm it imposes is not excessive in relation to the expected benefit.

³We should be careful, though, not to overstate the extent to which we should expect an account of subjective proportionality to be more action guiding than an account of objective proportionality. Just as we might be mistaken about facts concerning others (e.g., the threat they pose), so too might we make mistakes about our own evidence. Our own evidence is not always transparent to us. Moreover, even when our evidence is transparent, the correct theory of subjective proportionality might turn out to be too computationally demanding for agents to apply under certain conditions, as when police must make split-second decisions during violent encounters. For these reasons, the difference in action-guidingness between a theory of subjective proportionality and a theory of objective proportionality is more a difference in degree than in kind.

⁴Ibid., 255.

An initial problem for SIMPLE is that it ignores many important non-consequentialist considerations. On this account, for example, judgments of subjective proportionality are insensitive to facts about who is harmed, how harm is distributed, whether persons are harmed as means or as side-effects, or whether the harm is intended or merely foreseen. But SIMPLE can perhaps easily enough be modified to incorporate these considerations. We might incorporate them by adding a weighting factor into the mix:

SIMPLE WEIGHTED VIEW. A use of force is subjectively proportionate if and only if the expected *weighted* harm it imposes is not excessive in relation to the expected *weighted* harm it prevents.

Filling out such a theory would of course require specifying how this weighting works. It would need to specify, for example, how intended harm is weighted relative to merely foreseen harm. Tomlin argues that this is no easy task. But his primary worry with SIMPLE (and SIMPLE WEIGHTED) lies elsewhere.

What he sees to be the fundamental problem with the view is that it rests on what he calls the

Separation Assumption. Subjective proportionality assessments must consider all the potential harms that may result from a use of force, and then, separately, all of the potential benefits. The collective of the potential harms must then be compared with the collective of potential benefits.⁵

Tomlin makes a compelling case against this assumption. His objection appeals to cases like the following:

⁵Ibid., 257.

Variance. An attacker poses a threat of harm to you — a threat such that it would be objectively proportionate to impose on him up to 10 units of harm. You have two defensive options that will equally well avert his attack. Option one is (on your evidence) 50% likely to impose 9 units of harm on the attacker and 50% likely to impose 7 units of harm on the attacker. Option two is (on your evidence) 50% likely to impose 16 units of harm on the attacker and 50% likely to impose 0 units of harm on the attacker.

	State ₁	State ₂	Expected harm imposed
Probability	0.5	0.5	
Option ₁	9 units	7 units	8 units
Option ₂	16 units	0 units	8 units

SIMPLE does not discriminate between these options, since they impose the same expected harm (8 units) and provide the same expected benefit (they both succeed in fully averting the attack). But that’s clearly the wrong result; considerations of proportionality clearly favor the first option. There is no chance that you impose *objectively* disproportionate harm on option one, but there is a 50% chance that you impose objectively disproportionate on option two. Intuitively, that’s enough to at least "break the tie" between two option that are identical with respect to expected good and expected harm. Surely it should make a difference — at the subjective level — whether there is a possibility of doing what is objectively disproportionate.

The problem with SIMPLE — or any theory that makes the Separation Assumption — is that it tells us to calculate expected harm and expected benefit *before* engaging in any proportionality comparisons. The lesson of a case like Variance is that this gets things backwards: we need to engage in some proportionality comparisons *within* options before we start calculating any probability-weighted sums *across* options. And this is because

we need an account on which *subjective* proportionality is sensitive to the likelihood that an action will be *objectively* disproportionate.

2.2 Tomlin’s Metaproportionality Alternative

Tomlin develops an account that is thus sensitive. His proposal is to inject an additional measure — above and beyond expected harm and expected benefit — into the calculus. He proposes a measure of what he calls

Expected Objective Disproportionality (EOD). The *expected objective disproportionality* of an option is the probability-weighted sum of the harms — at each possible outcome — that are in excess of the amounts of harm that would be proportionate at that outcome.⁶

To illustrate, consider the case of Variance. Option one has an expected EOD of 0. This is because there is no chance of imposing disproportionate harm on that option. Option two, however, has a 50% chance of imposing 16 units of harm on the attacker. This would be disproportionate, and the amount of harm "in surplus" of the proportionality limit is 6 units. So option two would have a 50% chance of imposing 6 units of disproportionate harm and a 50% chance of imposing 0 units of disproportionate harm, giving us an EOD value of 3.

	State ₁	State ₂	EOD
Probability	0.5	0.5	
Option₁	proportionate	proportionate	0
Option₂	disproportionate by 6 units	proportionate	3

The appeal to such a value is attractive. We need some way to take account of the possibility of imposing objectively disproportionate harm. More

⁶Ibid., 277-278. Tomlin actually considers two different ways of defining this value, but focuses in his discussion on the formulation above.

precisely, we need it to come out that (i) all else equal, the lesser the *chance* of imposing objectively disproportionate harm the better, and (ii) all else equal, the lesser the *magnitude* of potentially objectively disproportionate harm the better. Tomlin's appeal to *expected objective disproportionality* is a natural way to capture both of these features.

But there remains the question what to do with this value. How exactly should this value interact with values like expected harm and expected benefit to determine which actions are subjectively proportionate?

Tomlin first considers two accounts on which expected objective disproportionality takes lexical priority over considerations of expected harm. On the first account, a use of force is subjectively proportionate just in case the EOD value is 0 — that is, just in case there is no possibility of that use of force being objectively disproportionate. On the second account, a use of force is subjectively proportionate just in case it *minimizes* EOD as compared to one's alternative defensive options.⁷

Tomlin rightly dispenses with such lexical views on the grounds that they are implausibly demanding. We can see this by considering, for example:

Great Odds. You must choose between (i) letting one-thousand people die to rising floodwaters, or (ii) taking a gamble in which there is a 99.99% chance that you will save the lives of the one-thousand and a 0.01% chance that you'll fail to save them and kill an innocent person besides.

Lexical views such as those above tell us that it is subjectively disproportionate to take the latter gamble. That's wildly implausible.

Tomlin considers two non-lexical alternatives. The first I'll call the

⁷Ibid., 278-279.

BOOSTING VIEW. A use of force is subjectively proportionate if and only if the [the expected (weighted) harm it imposes *plus* the expected objective disproportionality] is not excessive in relation to the expected benefit.

A view like this makes subjective proportionality sensitive to the possibility of imposing objectively disproportionate harm, while also permitting expected objective disproportionality to trade off against considerations of expected harm. For this reason, BOOSTING gets the right result in Great Odds. It tells us that the risky option is subjectively proportionate, since the expected harm it imposes is very small and the EOD value is very small, while the expected benefit is massive. The aggregate of the former values is not close to being excessive as compared to the latter value.

Tomlin, however, takes issue with BOOSTING. It involves, he argues, a problematic form of "double counting". Consider a case like:

Two Attackers. Aggressor₁ threatens harm to Alice; Aggressor₂ threatens equally serious harm to Bob. It is proportionate to impose up to 10 units of harm on either aggressor. But you can only prevent one attack or the other. You can defend Alice only by harming Aggressor₁ to degree 10. You can defend Bob only by harming Aggressor₂ to degree 12.

BOOSTING counts the good of defending Alice only *once*: it counts towards the subjective proportionality of harming Aggressor₁ by contributing to the amount of expected benefit that action. But BOOSTING counts the good of defending Bob *twice*: it contributes to the amount of expected benefit, but also counts against the expected objective disproportionality of that action. Tomlin finds this an implausible result. Why should Bob's interests make a greater moral difference than Alice's?

Tomlin suggests an alternative approach. Objective proportionality involves a comparison between harm and benefit. It's natural, then, to think that subjective proportionality would involve a parallel comparison between *expected* harms imposed and *expected* benefit. But this is where we'd be mistaken, says Tomlin. We're looking at the wrong comparators. Subjective proportionality involves, he says, not the comparison of expected harm with expected benefit, but a comparison of the expectation of *objectively disproportionate harm* to the expectation of *benefit achieved by objectively proportionate means*.⁸ There are different ways of spelling out the details. For the purpose of illustrating the view, here's one precisification:

META VIEW. A use of force is subjectively proportionate if and only if the expected objective disproportionality (EOD) of that action is greater than its expected objective proportionality (EOP). An action's EOP is the probability-weighted sum of the surplus of benefit over harm for each of the objectively proportionate outcomes of that action.⁹

⁸"Proportionality, in the simplest terms, is about comparing bad things with good things. Here, the bad things are the objectively disproportionate potential outcomes ... Conversely, the good things are the objectively proportionate potential outcomes. What we need to know is whether it is worth risking the bad potential outcomes for the sake of the good potential outcomes. So, our focus in looking for the good against which to balance these bad outcomes should be on the objectively proportionate potential outcomes" (Ibid., 281).

⁹Tomlin actually takes EOP to be a bit more complicated than this. He argues that, when calculating EOP, we should ignore harms that are imposed on persons who are liable to suffer those harms (i.e., who are not wronged by those harms). An important reason for Tomlin to make this move is that such a move is needed for META to get the right result in a case like Variance: if the harms suffered by the attacker count towards EOP, then the difference between EOP and EOD is the same on option₂ as on option₁. But getting a difference in subjective proportionality between two options in a case like Variance was the whole motivation for giving up the Simple View in the first place!

I worry, however, about the plausibility of entirely ignoring harms to liable parties. As numerous theorists have noted, we should perhaps discount harms suffered by liable parties, but it doesn't seem we should count them for naught, especially when we consider cases involving "minimally-responsible" attackers. See, for example, Seth Lazar, "Necessity in Self-Defense and War," *Philosophy & Public Affairs* 40 (2012): 12;

The core idea is that only certain benefits count towards subjective proportionality — namely, the benefits that obtain on objectively proportionate outcomes.¹⁰ To illustrate, consider:

Rail Bombing. A military commander on the just side of a war must decide whether to bomb an important enemy rail line. On his evidence there are three equiprobable outcomes. The best outcome is that the bomb hits its target and detonates below the surface, minimizing collateral damage: 10 units of benefit for 1 unit of harm. The second outcome is that the bomb hits its target but detonates without penetration, causing more (but still proportionate) collateral damage: 10 units of benefit for 7 units of harm. The third outcome is that the bomb misses the rail line, causing disproportionate harm: 0 units of benefit for 9 units of harm.

	Subsurface hit	Surface hit	Miss
Probability	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
Bombing harm	1	7	9
Bombing benefit	10	10	0

META tells us to calculate the expected objective disproportionality (EOD) of dropping the bomb and compare this value to the expected objective proportionality (EOP) of dropping the bomb. (To simplify things, let us imagine that 1 unit of benefit justifies up to 1 unit of harm — i.e., let's assume a proportionality ratio of 1:1.) To calculate EOD, we look

and McMahan, "The Limits of Self-Defense." A more plausible move for Tomlin, I think, would be to simply discount harms that fall on liable parties when counting these harms against the EOP of an option, without ignoring those harms entirely. This would allow META to get the right result in Variance while also acknowledging that we have reasons to not impose harm even on liable parties.

¹⁰"Any coupling of harm and good that is objectively disproportionate counts against an action, and if we are to risk such outcomes, we must justify them in terms of the objectively proportionate potential outcomes" ("Subjective Proportionality," 282).

only at harm in outcomes that are objectively disproportionate. Only the "miss" outcome is objectively disproportionate, and by 9 units of harm. Weighting for probability, this gives us an EOD value of 3.

To calculate EOP, we look only at benefits in outcomes that are objectively proportionate. Both outcomes where the bomb hits the rail line are objectively proportionate. The "benefit surplus" on the outcome where the bomb detonates on the surface is 3. The benefit surplus on the outcome where the bomb detonates subsurface is 9. Weighting for probability, we get an EOP of 4.

Since $EOP > EOD$, dropping the bomb is subjectively proportionate, according to META.

3 Rethinking the Simple View

I think Tomlin is surely right that we must reject the Separation Assumption. Cases like Variance clearly demand that we make some intra-option proportionality comparisons before aggregating across outcomes. In what follows, however, I want to motivate an alternative (and simpler) way of thinking about what role these comparisons should play in our decision-making under uncertainty.

My proposal is founded on a worry that Tomlin's appeal to a value like *expected objective proportionality* glosses over — in a problematic way — important differences between what theorists call "narrow" and "wide" proportionality.

Let's start by getting clear on that distinction, focusing first on simple cases where outcomes are certain. *Narrow proportionality* concerns the way in which a certain relation of proportionality fixes a limit on the rights that are forfeited by an aggressor. To illustrate, compare:

Attempted Murder. Aggressor₁ attempts to kill you. He will

succeed unless you kill him first.

Attempted Pinch. Aggressor₂ attempts to give you a pinch on the arm. He will succeed unless you kill him first.

You're permitted to kill the would-be murderer but not the would-be pincher. The would-be murderer has forfeited his right to not be killed; not so for the would-be pincher. This is because an attacker forfeits only his rights against defensive harms that are *proportionate* to the harm it prevents.

So the first place at which "proportionality" is relevant to the morality of harm is this: a relation of proportionality constrains the amount of harm that may be imposed on someone *without infringing their rights*. The relevant proportionality relation is what we'll call (following the literature) the *narrow proportionality relation*.¹¹

But there's another proportionality relation — the *wide proportionality relation* — that is relevant to the morality of harm. To say that harm is "widely" proportionate is not to focus on the relationship between the harm imposed on the attacker and the harm averted thereby. It is rather to say something about the balance of one's *all-things-considered* reasons for and against the use of harm. For defensive or punitive harm to be widely proportionate is for the reasons in favor of imposing that harm to outweigh one's reasons against imposing that harm.

It's very important that we not conflate these two forms of proportionality. They're importantly different from one another, in at least two ways. First, they involve different *relata*. Narrow proportionality involves a comparison between two amounts of harm — the harm imposed on the aggressor and the harm prevented thereby.¹² Wide proportionality,

¹¹See Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009), 22-24 for an early formulation of the narrow/wide proportionality distinction.

¹²Though there are different views about what this comparison involves, and about

by contrast, involves a contrast between two bundles of reasons. Second, these two relations play different roles. The narrow proportionality relation functions as a constraint on a particular individual's liability to harm. The wide proportionality relation functions, not as a constraint on anyone's liability, but rather as a constraint on all-things-considered permissibility. To say that defensive harm is narrowly disproportionate implies that the harm *wrongs* someone, but it doesn't imply that the harm is impermissible. To say that defensive harm is widely disproportionate implies that the harm is impermissible, but it doesn't imply that it wrongs anyone.

Importantly, these differences mean that defensive or punitive harm can be narrowly proportionate without being widely proportionate, and vice versa. Compare, for instance:

Small Grenade. An aggressor attempts to kill you. He will succeed unless you throw a grenade at him. The grenade will kill the attacker, and also cause each of ten innocent bystanders to suffer a mild scratch.

Big Grenade. As in Small Grenade, except that the grenade will *kill* the ten bystanders.

In each case, the harm is narrowly proportionate: the harm imposed on the attacker is proportionate to the harm averted thereby. In each case, the attacker is liable to be killed, and the bystanders are not liable to any harm. And yet only in the first case is it permissible to throw the grenade. This is because only in the first case is the defensive harm widely proportionate: only in the first case are the reasons in favor of throwing the grenade weighty enough to outweigh the reasons against throwing the grenade.

whether these are exactly the right comparators. See, for example, Clark, "Demands," 133; McMahan, "Culpability," 766; Quong, *Defensive Force*, 96-123; Tadros, *Ends*, 347; and Uniacke, "Proportionality," 253-272.

Likewise, harm can be widely proportionate without being narrowly proportionate. Someone attempts to pinch you; you can prevent the pinch only by breaking their legs. It would be narrowly disproportionate to break their legs. But suppose that breaking their legs would deter a different aggressor from bombing a crowded office building. It would then, I think, be widely proportionate to break their legs.¹³

Tomlin is sensitive to the distinction between narrow and wide proportionality. And indeed he purports to develop a theory that incorporates both forms of proportionality into decision-making under uncertainty in a uniform way: "In this paper," he says, "I will often focus on narrow proportionality cases, but everything I say is, unless I indicate otherwise, supposed to apply to wide and narrow proportionality alike." What this means, I take it, is that his proposed value of *expected objective disproportionality* is one that measures both narrow and wide objective disproportionality. An outcome that would be merely widely disproportionate will contribute to expected objective disproportionality just as would an outcome that would be merely narrowly disproportionate.

I think this a crucial mistake. My concern has to do with the fact that narrow and wide proportionality are grounded in different moral concerns. Let me explain.

It matters that we avoid doing what is narrowly disproportionate because it matters that we not infringe others' rights. It matters that we avoid doing what is widely disproportionate because it matters that we do what we have most reason to do. These are different considerations that seem to behave differently under conditions of uncertainty.

¹³Narrow and wide proportionality are importantly distinct, but not entirely independent. Facts about narrow proportionality are among the grounds of facts about wide proportionality. Defensive harm that is narrowly disproportionate is (for this reason) harm that is *rights infringing*. We have strong reasons not to infringe rights, and thus the fact that some harm would be narrowly disproportionate makes it much harder for that harm to satisfy the wide proportionality constraint.

To see why, let's start by recalling the case of Variance, where the attacker makes himself liable to suffer up to (but no more than) 10 units of harm.

	State ₁	State ₂	Expected harm imposed
Probability	0.5	0.5	
Option₁	9 units	7 units	8 units
Option₂	16 units	0 units	8 units

Earlier we followed Tomlin in noting that it matters that option two has the chance of imposing narrowly disproportionate harm. But here's why that matters in this case: it's because the chance of imposing narrowly disproportionate harm means the chance of *infringing someone's rights*, and it matters greatly that we not infringe others' rights. That gives us reasons not to impose that harm above and beyond the severity of the harm itself.

But now, by contrast, consider an analogous case where it is only wide disproportionality that is risked:

Wide Variance. A fire engulfs an office building. You can prevent serious burns to five people by redirecting the flames such that Alice (a bystander) suffers serious burns. Or you can save the five by redirecting the flames such that there is a 50% chance that no one will be harmed and a 50% chance that Alice and Bob both suffer serious burns.

	State ₁	State ₂	EOD	Expected Harm
Probability	0.5	0.5		
Option₁	1 (burned)		0	1
Option₂	0	2	1	1

Let's stipulate that it is widely proportionate to burn one, but no more than one, person as a side-effect of saving five. (If those numbers seem

off, adjust to taste.) Given those stipulations, only your second option has any chance of imposing objectively widely-disproportionate harm. On Tomlin's approach, then, the second option will have a greater EOD value.

Intuitively, however, Wide Variance is not like Variance. In Variance, it is plainly worse to choose the second option (the option that risks narrowly disproportionate harm). In Wide Variance, by contrast, it is not at all obvious that you should choose the second option (risking widely disproportionate harm). On the contrary: I think your two options are morally on par in Wide Variance. Here's a justification for that idea. In a case like Variance, the chance of narrowly disproportionate harm matters because it matters that we not infringe rights. The expectation that we wrong others must be accounted for above and beyond the expectation of harm. But in a case like Wide Variance, you risk infringing someone's rights on either option. Your first option is guaranteed to (permissibly) infringe one person's rights; your second option will either infringe no one's rights or infringe two person's rights. The expected severity of rights-infringement is identical on both options. One way to see this is to consider the position of Alice and Bob. They seem to have equally strong, opposing complaints. If you choose option₁, Alice could complain that you could have chosen a different option that would have given her a 50% better chance of not being wrongfully burned. If you choose option₂, Bob could complain that you could have chosen a different option that would have given him a 50% better chance of not being wrongfully burned. In as much as our concern is with the infringement of rights, there seems no reason to privilege one option over the other. This is what makes Wide Variance so different from Variance.

Now it's of course true that, in Wide Variance, only option₂ has a chance of imposing widely objectively-disproportionate harm. But to say that an action is widely disproportionate is just to say that the balance of reasons does not favor that option. What's risked by choosing option₂, then, is

the possibility that the reasons in favor of imposing harm will fail to outweigh the reasons against. But if the *balance of reasons* is what we're worried about, then wouldn't this risk be offset by the potential upside of taking the gamble that is option₂? That option risks killing one person too many, but it also risks killing no one. Because your reasons not to kill Alice are just as weighty as your reasons not to kill Bob (and vice versa), option₁ and option₂ will have the same expected weight of reasons in their favor. It's hard to see, then, why the moral concern that grounds the wide proportionality constraint should lead us to prefer the one option over the other.

Let me suggest an alternative formulation of subjective proportionality in light of these observations about the differences between narrow and wide proportionality. We should, I think, maintain the basic structure of the [SIMPLE VIEW](#), but replace the focus on expected harm with a focus on expected reasons.

SIMPLE REASONS. A use of force is subjectively proportionate if and only if the expected weight of reasons in favor of that use of force is greater than the expected weight of reasons against that use of force.¹⁴

The appeal to reasons best captures, I think, the moral considerations that underlie our concern with both narrow and wide proportionality at the objective level. It is important that our actions are narrowly proportionate because it is important that we not infringe rights; we have very weighty reason not to infringe rights, apart from the harm that may be caused thereby. **SIMPLE REASONS** captures this concern by giving weight to the possibility that an action will infringe someone's rights. (For this reason,

¹⁴For **SIMPLE REASONS** to be action-guiding, of course, the account needs to be supplemented with information about the weight of our reasons not to infringe rights, as compared to the weight of our reasons not to harm and the weight of our reasons not to allow harm.

SIMPLE REASONS departs from the original SIMPLE account, in that the former is incompatible with the [Separation Assumption](#). This is because harm at a possible outcome is wrongful if it is narrowly disproportionate at that outcome, and whether harm is narrowly disproportionate at an outcome depends on proportionality comparisons that must be made *within* that outcome.)

The moral considerations that underlie our concern with wide proportionality are different. It matters that our actions are widely proportionate because it is permissible to perform an action only if there is more reason to perform the action than not. SIMPLE REASONS captures this concern by treating subjective proportionality as a subjective analog of wide proportionality. Where an action is objectively, widely proportionate just in case there is more reason to perform the action than not, an action is subjectively proportionate just in case there is more expected reason to perform the action than expected reason not to. On this approach, there is no useful distinction to be made between "subjective narrow proportionality" and "subjective wide proportionality". Subjective proportionality is only ever "wide," but it takes account of the considerations for which we care about narrow proportionality, by taking account of the fact that we have stronger reason to avoid narrowly disproportionate harm than narrowly proportionate harm of equal severity.

That's a crucial difference between SIMPLE REASONS and Tomlin's approach. Where Tomlin's appeal to *expected objective disproportionality* makes subjective proportionality sensitive to the expectation of both narrow and wide objective proportionality, SIMPLE REASONS tells us that subjective proportionality is sensitive only to the expectation of narrow objective proportionality. The moral value that undergirds the importance of wide proportionality is accommodated without an appeal to the expectation of wide disproportionality.

Notice that SIMPLE REASONS is a view that comes very close to a version of

what we called (above) the **SIMPLE WEIGHTED VIEW**. It's a version of that view where the weighting function gives extra weight to rights-infringing harm over non-rights-infringing harm. Again, though, it's important to notice that this feature makes the view incompatible with the Separation Assumption. This is because we can only assess whether harm wrongs someone in virtue of being narrowly disproportionate by making proportionality comparisons *at an outcome*. Intra-outcome proportionality judgments are relevant to determining whether someone's rights have been infringed, and are thus relevant to the quantity of "weighted" expected harm.

Now it might be worried that **SIMPLE REASONS** is subject to Tomlin's "double counting" worry about the **BOOSTING** account.¹⁵ Recall the case of:

Two Attackers. You can defend Alice from Aggressor₁ only by harming the latter to degree 10. You can defend Bob from Aggressor₂ only by harming the latter to degree 12. You can only defend one of Alice or Bob, and it would be narrowly disproportionate to impose more than 10 units of harm on either attacker.

Tomlin's worry for **BOOSTING** was that it unfairly counts the good of defending Alice once but the good of defending Bob twice. Preventing harm to Alice counts towards the subjective proportionality of harming Aggressor₁ only by contributing to the amount of expected benefit of

¹⁵**SIMPLE REASONS** and **BOOSTING** have important extensional differences. First, where an outcome involve narrowly disproportionate harm, **BOOSTING** merely double counts the disproportionate harm. **SIMPLE REASONS** gives considerably more weight to narrowly disproportionate harm than this — at least on the very plausible assumption that it is *more than twice as bad* to impose n-units of rights-infringing harm on someone than to impose n-units of non-rights-infringing harm on someone. But while **SIMPLE REASONS** takes the possibility of narrow proportionality more seriously than **BOOSTING**, it takes the possibility of wide proportionality less seriously. Where an outcome involves only widely disproportionate harm, **BOOSTING** again double counts the widely disproportionate harm. **SIMPLE REASONS**, by contrast, gives no extra weight to this harm.

doing so. Preventing harm to Bob counts twice: it contributes to the amount of expected benefit but *also* counts against the expected objective disproportionality of that action.

Should we be similarly worried about SIMPLE REASONS? I don't think so. On this account, everyone's interests contribute in the same ways. The value of preventing harm to Bob counts at two places. First, it contributes to fixing the amount of harm to which Aggressor₂ is liable. Second, it contributes to the expected benefit of killing that aggressor. But the value of preventing harm to Alice likewise counts at two places. It contributes to fixing the amount of harm to which Aggressor₁ is liable and to the expected benefit of killing that aggressor. There are no worries about inequality here. And, moreover, I can't see what would otherwise be objectionable about Alice and Bob's interests playing this dual role of both (i) fixing an attacker's liability and (ii) contributing to the expected benefit of harming the attacker.

4 Conclusion

The [Separation Assumption](#) is false. But it's important to notice why it's false. The crucial problem is that it fails to accommodate the moral concern that underlies the importance of narrow proportionality. We have especially strong reasons not to impose rights-infringing harm. This means that our decision-making under uncertainty must give special weight to rights-infringing harms. But among the determinants of whether a possible harm would be rights-infringing are facts about whether that harm would be narrowly disproportionate. Imposing narrowly disproportionate harm on someone is one way of infringing someone's rights. So we have to make proportionality judgments *within* a possible outcome before aggregating harms (or reasons) *across* outcomes, against the Separation Assumption.

We must be careful, however, not to blur the important differences between narrow and wide proportionality. The former is grounded in the

importance of not infringing rights. The latter is grounded in the importance of having sufficient all-things-considered reasons for action. We need to make intra-outcome narrow proportionality judgments to incorporate the former ground into our decision-making. But there's no need to make intra-outcome wide proportionality judgments in order to give the latter concern its due weight. On the contrary, the best way to give this concern its due is by thinking of subjective proportionality as an analog of objective, wide proportionality. Where an action is widely proportionate just in case the weight of reasons in favor of the action are greater than the weight of reasons against, an action is subjectively proportionate just in case the weight of expected reasons in favor are greater than the weight of expected reasons against.