

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/104423>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Subjective Proportionality*

Patrick Tomlin

forthcoming in Ethics

Abstract: Philosophers writing about proportionality in self-defence and war will often assume that defensive agents have full knowledge about the threat that they face and the defensive options available to them. But no actual defensive agents possess this kind of knowledge. How, then, should we make proportionality decisions under uncertainty? The natural answer is that we should move from comparing the harm we will do with the good we will achieve to comparing expected harm with expected good. I argue that this simple calculation is flawed, and begin to develop a more sophisticated account of ‘subjective proportionality’.

I INTRODUCTION

Philosophers who write about proportionality in self-defence and war will often set things up in this way: Attacker threatens Victim with harm. Victim knows that if she does nothing, she will suffer x units of harm. Victim also knows that if she does Action A, she will inflict y units of harm on to Attacker, and that this will produce outcome z (e.g., will fully prevent the attack). In other words, perfect knowledge is assumed about the attack and the defensive options available to the defensive agent. This is unproblematic insofar as we are interested in what Thomas Hurka calls ‘objective proportionality’¹, where we weigh the actual harm caused against the actual good it

* I have benefitted from discussion with, and comments from, Helen Frowe, Alex Gregory Adil Haque Robert Jubb, Seth Lazar, Kieran Oberman, Jonathan Parry, David Plunkett, and Steve Woodside. The paper was presented at the Stockholm Centre for the Ethics of War and Peace, Normative Orders at Goethe Universität, Frankfurt, and the Universities of Oxford, Southampton, Glasgow, Pennsylvania, and Reading. Once submitted, the paper was much improved by several sets of comments and clear guidance from Cécile Fabre, David Wasserman, and an anonymous *Ethics* referee. I owe special thanks to two people. The first is Jeff McMahan. It was at his urging that I first began to write up some inchoate thoughts on proportionality and risk, and he has offered help and guidance at many points since. The second is Victor Tadros, for many discussions on these topics. While writing the paper, I received support from Goethe Universität’s *Justitia Amplificata* and *Forschungskolleg Humanwissenschaften*.

¹ Thomas Hurka, ‘Proportionality in the Morality of War,’ *Philosophy & Public Affairs* 33 (2005): 34-66, at 38. Alternatively, we can call this ‘fact-relative proportionality’. For the idea of fact-relative proportionality, see Jeff McMahan, ‘Proportionality and Time,’ *Ethics* 125 (2015): 1-25, at 11, drawing on Derek Parfit, *On What Matters*, Vol 1 (Oxford: Oxford University Press, 2011), ch. 5.

achieves.² But few, if any, defensive agents (be they private individuals, states, or soldiers) will ever face a violent situation in which they possess perfect knowledge of x, y, and z. Therefore, how we should make proportionality calculations under conditions of empirical uncertainty is an important moral and practical question in the ethics of war and self-defence.³ We can call this the question of ‘subjective proportionality’.⁴

Developing an account of subjective proportionality is important for 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. Second, since it takes into account the agent’s epistemic position, such an account will be essential in developing a standard on the basis of which we can judge behaviour, and so attribute praise, blame, and, potentially, punishment.⁵

Few philosophers have written explicitly on this topic. But there does seem an obvious way to move from objective proportionality to subjective proportionality. Since, put

² It is important to note that I claim that assuming away uncertainty about harm and good is unproblematic for objective (or fact-relative) proportionality, not objective (or fact-relative) permissibility. I leave it open here whether risk can be relevant to fact-relative permissibility. For example, I take no stand on Seth Lazar’s claim that ‘risky killings’ are ‘objectively worse’ (and therefore sometimes objectively impermissible, in virtue of the ex ante risks). See: Seth Lazar, ‘Risky Killing: How Risks Worsen Violations of Objective Rights,’ *Journal of Moral Philosophy* (forthcoming); *Sparing Civilians* (Oxford: Oxford University Press, 2015), ch. 4; ‘Risky Killing and the Ethics of War,’ *Ethics* 126 (2015): 91-117. I am grateful to Seth Lazar, Steve Woodside, and an *Ethics* referee for useful comments here.

³ As will become clear, I use ‘uncertainty’ in a non-technical sense here, since I presume that we can attach probabilities to various outcomes.

⁴ Following Derek Parfit’s distinction between evidence- and belief-relative permissibility, Jeff McMahan suggests that what Hurka and I call ‘subjective proportionality’ can usefully be divided into evidence- and belief-relative standards of proportionality. The considerations that I argue for here should be taken into account by both evidence- and belief-relative proportionality, since they concern how to take intentions and probabilities into account, and the probabilities in question could be evidence- or belief-derived. For simplicity, I will assume here that defensive agents believe and act on the best evidence available, thereby collapsing McMahan’s distinction. See: McMahan, ‘Proportionality and Time,’ 11.

⁵ There are two different direct roles that subjective proportionality might play in assigning blame and punishment. It might be a necessary condition of justification. Or it might be an excuse for having produced an objectively disproportionate, and thus unjustified, outcome. However, subjective proportionality’s role in a blameworthiness or legal standard may be indirect, since the right theory of subjective proportionality may be too complex to expect, for example, soldiers to use it in the heat of battle. If that is the case, heuristics, rules of thumb, and realistic laws would need to be developed.

simply, proportionality involves comparing harm caused with good produced⁶, when moving from objective to subjective proportionality it is natural to move from comparing actual harm with actual good to comparing expected harm with expected good. (These ‘expected’ values are calculated by looking at each possible outcome, multiplying it by the probability that it will eventuate, and then summing these figures). This is how Thomas Hurka suggests we do things in his important and influential essay ‘Proportionality in the Morality of War’, going so far as to suggest that it is the only plausible way for proportionality calculations to take uncertainty into account. Hurka claims that when assessing a prospective war’s proportionality if the war has only some small probability of achieving relevant goods ... then its expected harm is excessive compared to its expected good. If it takes account of probabilities in this way, as on any plausible view it must, the *ad bellum* proportionality condition incorporates hope-of-success considerations, and it can also incorporate last-resort considerations.⁷

From this, we can draw two key claims.

- (1) On any plausible view, subjective proportionality assessments must take probabilities into account.
(1) is surely correct. I will not question it here.
- (2) On any plausible view, probabilities should be taken into account by checking whether expected harm is excessive when compared to expected good. Call this *Hurka’s Test*.

Hurka’s Test can give rise to this Simple View of subjective proportionality:

⁶ Throughout I will assume that proportionality involves measuring harms against goods. This is controversial, but nothing I say here hangs on this particular comparison. Others are free to substitute their own metrics. For alternative metrics, see Jonathan Quong, ‘Proportionality, Liability, and Defensive Harm,’ *Philosophy & Public Affairs* 43 (2015): 144-173; F.M. Kamm, *Ethics for Enemies: Terror, Torture, and War* (Oxford: Oxford University Press, 2011), 133-134.

⁷ Thomas Hurka, ‘Proportionality in the Morality of War,’ 37.

Simple View: Hurka's Test exhausts subjective proportionality. A defensive action (or course of action) is subjectively proportionate if and only if its expected harm is not excessive compared to its expected good.⁸

I have two main aims in this paper. The first is to show that the Simple View is false. I reject the Simple View because – even if we complicate it a little, as we will – it ultimately rests on this assumption, which I will argue is fundamentally flawed:

The Separation Assumption: subjective proportionality assessments must consider all of the potential harms that may occur, and then, separately, all of the potential goods that may result. The collective of potential harms must then be compared with the collective of potential goods.

Put simply, I reject this Separation Assumption because it separates the harms from the goods, and this rules out taking account of what *pairings* of harm and good may come into existence together. In particular, the Separation Assumption therefore rules out taking into account whether or not a defensive action will be, or is likely to be, objectively disproportionate. I will argue that subjective proportionality must take this into account.

My second aim is to begin to develop a more sophisticated, and more plausible, account of subjective proportionality. In doing so, I will come to reject not only the Simple View, but any role whatsoever for Hurka's Test: the direct comparison between expected harm and expected good is, in my view, irrelevant to subjective proportionality.

The article proceeds as follows. In the next section, I introduce some key concepts and distinctions from the literature on proportionality in self-defence and war. I also introduce some different ways in which defensive agents may be uncertain when faced

⁸ The Simple View is suggested by Hurka's text, but cannot be attributed to him, since he says that proportionality must 'take account of probabilities', and that the way to do so is through Hurka's Test. He does not, however, explicitly state that this test is the only test (though he does not mention any others). I will later reject any role for Hurka's Test. I am grateful to Hurka for some useful feedback and for help with clarifying his view (or at least his view in this paper – I do not claim it is his current view).

with a threat. In Section III I complicate the Simple View a little, so that it is able to take account of some of the distinctions and concerns introduced in Section II, and other non-consequentialist concerns. I identify the key assumption – the Separation Assumption – which underpins any version of the Simple View.

I then attack this assumption on two bases. In Section IV I show that this assumption cannot properly accommodate the moral significance of intentional harm. Many philosophers working on self-defence and war consider intentional harm to be morally distinct from merely foreseen harm. I consider a variety of ways that one could try to accommodate the significance of intentional harm within the Simple View, and all are found wanting. I argue instead that it matters for proportionality whether or not the defensive agent's intended harm will lead to objectively proportionate outcomes, and that this means taking into account the pairing of the intended harm and the good it will produce. I delineate two different ways in which we could take account of this pairing, but both require us to reject the Separation Assumption. In Section V, I further examine this idea by focusing on what intended harm must be measured against. It matters, as I argue in Section IV, whether or not intended harm 'leads to' objectively proportionate outcomes. But, as I show in Section V, this claim is ambiguous between several readings. An intended harm may lead to a plurality of possible good outcomes, and some of these may be objectively proportionate, and others objectively disproportionate. I ask whether it matters that the collective of potential good outcomes that lie causally downstream from the intended harm are, taken together, acceptable, or whether we must justify the intended harm as proportionate to the good it is intended to produce. I explain why I think intended harm must be proportionate to the intended good.

In Section VI I present my second objection to the Separation Assumption – it ignores whether an action risks objectively disproportionate outcomes, how likely such outcomes are, and how seriously objectively disproportionate they might be. We can do so, I argue, if we take account of a measure I call Expected Objective Disproportionality. Section VII examines what role this measure might play in a theory of subjective proportionality. Section VIII concludes.

Through this discussion, the beginnings of a more plausible, and more complex, positive account of subjective proportionality will emerge. There are some elements of this account in which I have greater confidence than others. I will be clear about which elements I am unsure of as I proceed. Often, my aim will be less to defend a particular view, than to articulate a difficult issue for the emerging account, and to present the candidate views one might hold. I hope to show, at minimum, that not all plausible accounts of subjective proportionality involve the comparison between expected harm and expected good.

II PROPORTIONALITY, DEFENSIVE HARM, AND UNCERTAINTY

Before proceeding to my critique of the Simple View, it will be useful to briefly outline the role of proportionality in justifying defensive harm, some different kinds of proportionality that have been delineated in the literature, and how they interact with some other concepts in justifying defensive harm. I will also identify some different ways in which agents may be uncertain when faced with a threat.

Proportionality, as we understand it here, is a limit on violence. In punishment theory, proportionality is sometimes taken to be a just aim – both less-than-proportionate punishments and more-than-proportionate punishments are seen as unjust. But proportionality in self-defence and war is not a just aim, it is a limit on the pursuit of just aims: proportionality condemns more-than-proportionate harms, but is silent on harms that fall below the limit that it establishes. For a given level of good to be achieved, proportionality establishes a maximum amount of proportionate harm. For a given level of harm, proportionality established a minimum amount of good that must be achieved.

Proportionality is not the only limiting concept in justifying defensive force. If a plurality of possible defensive actions are proportionate, the principle of necessity

states that (all else equal) we should pick the option that minimizes harm.⁹

Proportionality and necessity thus act in concert, as independent limits on defensive force: proportionality demands that defensive acts do enough good, given the harm they will inflict; necessity demands that defensive acts do no more harm than is necessary to achieve a good outcome.

The literature on proportionality, like work on the ethics of self-defence and war more generally, has become much more sophisticated in recent years. Jeff McMahan has identified a key distinction within the concept of proportionality, between what he terms ‘narrow proportionality’ and ‘wide proportionality’.¹⁰ Narrow proportionality concerns defensive harm to those who are culpable (or perhaps merely responsible) for a threat, and are thus potentially liable to the defensive harm. When an agent is liable to harm, she has no right against it, and thus is not wronged by the harm. Wide proportionality concerns harm to innocent third parties, which require a ‘lesser-evil’ justification. Unlike liability-justified harms, lesser-evil harms infringe rights, but provided enough good is done, these infringements can be widely proportionate, and thus potentially permissible. In this paper 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.

Alongside McMahan’s narrow-wide distinction, another important distinction for us here is that between objective and subjective proportionality. Objective proportionality is used to evaluate outcomes – particular pairings of harm and good that come into the world together. Subjective proportionality concerns whether an act is proportionate given the evidence available to the agent, or the beliefs of the agent.¹¹ In examples in which there is only one possible outcome from a defensive action, the two forms of proportionality map precisely onto one another – if the act is certain to

⁹ This is a simplistic statement of the necessity principle, but it will do for our purposes here. For complexities, see: Seth Lazar, ‘Necessity in Self-Defense and War,’ *Philosophy & Public Affairs* 40 (2012): 3-43; and Jeff McMahan, ‘Proportionality and Necessity in *Jus in Bello*’ in Helen Frowe and Seth Lazar, eds., *The Oxford Handbook of the Ethics of War* (New York: Oxford University Press, 2017).

¹⁰ Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009), 19-29

¹¹ As previously stipulated, for the purposes of this essay I assume that agents believe and act upon the best evidence available to them.

be objectively disproportionate, it is also subjectively disproportionate, and the agent ought not to perform it. However, under uncertainty these two kinds of proportionality can come apart. Consider a case where, in order to avert a minor harm to myself (such as a scratch) I perform a defensive act which I reasonably believe will cause minor harm to an attacker. However, things go wrong, and I end up unintentionally killing him. We need moral language to describe the act from both the subjective and objective perspective. The act turned out to be disproportionate – it did not do enough good to justify that level of harm, and had I known I would cause that level of harm, I would not have been permitted to proceed. So, it is objectively disproportionate. However, it was also subjectively proportionate – the evidence suggested I would do very little harm.

There are two key ‘inputs’ to an objective proportionality calculation – the harm the defensive agent does, and the good it achieves. When we turn our focus to subjective proportionality, the defensive agent can (and ordinarily will) face uncertainty about both. Uncertainty about the amount of harm the defensive agent will do is easy enough to understand: if Victim attempts to punch Attacker, she might give him a black eye, but there’s also a small chance she will kill him, miss altogether, or accidentally harm Bystander. Let’s call this Harm Uncertainty. Harm Uncertainty can concern whether a given action will succeed in harming at all, the extent of any harm, and who will suffer it.¹² There can also be uncertainty about the moral status of those who will be harmed – I can be sure I will harm Bob, but unsure as to whether he is innocent or responsible, and thus whether we are dealing with narrow or wide proportionality.¹³

¹² Interestingly, in his ‘Justifying Harm’ David Rodin appears to be outlining an approach to subjective proportionality, as he takes some forms of uncertainty into account – in particular what I call Threat Uncertainty and Outcome Uncertainty. And yet when Rodin comes to consider the defensive harm, he does not take account of Harm Uncertainty (see 92). (However, see 97 where he refers to ‘the magnitude and expected probability of the threatened and defensive harms’, and 105). This omission means that Rodin’s account is silent on what I show here is a very complex issue: how to combine Harm Uncertainty and a concern for intentional harm. See: David Rodin, ‘Justifying Harm,’ *Ethics* 122 (2011): 74-110,

¹³ For comments on this latter kind of uncertainty in a war context, see Seth Lazar, ‘The Responsibility Dilemma for Killing in War: a review essay,’ *Philosophy & Public Affairs* 38 (2010): 180-213.

Uncertainty about the good that the defensive agent will do, however, is more complicated, and can, for our purposes, be broken down usefully into two subtypes, which, temporally-speaking, fall either side of the harm that the defensive agent inflicts, or attempts to inflict. The first, which precedes the defensive harm, is Threat Uncertainty. This is uncertainty about the level of the threat Victim faces. Threat Uncertainty frames the entire choice-making situation that the defensive agent faces: since she doesn't know how much harm she will suffer if she allows the threat to eventuate, she does not know how much harm she will prevent, and thus how much good she will do, by fully averting the threat.

The second kind of uncertainty about the good a defensive action will achieve is Outcome Uncertainty. Outcome Uncertainty is like Threat Uncertainty, in that it is uncertainty concerning how much good a defensive action will do. But it is also like Harm Uncertainty, in that it is uncertainty about the defensive agent's own actions and their potential consequences. Outcome uncertainty concerns the effectiveness of a defensive harming action (and therefore lies temporally downstream from Harm Uncertainty). For example, the defensive agent may be sure that she faces a given threat, and certain that she can punch the attacker, giving him a black eye, but uncertain as to whether this will deter him, make no difference, or anger him, thereby exacerbating the threat.

Whilst the Hurka quotation with which we began specifically references ad bellum proportionality, the present essay also concerns both in bello proportionality and proportionality in self-defence, and will draw upon self-defence examples. While the precise relationship between individual self-defence and the ethics of war is controversial, most philosophers believe that our thinking about one can rightly inform our thinking about the other.¹⁴ There are also controversial questions about whether ad bellum and in bello proportionality must be assessed in the same or

¹⁴ This includes Hurka ('Proportionality in the Morality of War,' 38-39). For an overview of the debate concerning the relationship between war and self-defence, see Helen Frowe, *The Ethics of War and Peace* (Abingdon: Routledge, 2011), ch. 2.

different ways.¹⁵ For the purposes of this essay, I am happy to allow that the ‘good’ in in bello proportionality may differ markedly from the relevant goods in ad bellum or self-defence proportionality calculations. I try to articulate, here, a framework, and different accounts of the relevant harms and goods can be plugged into that framework.

III NON-CONSEQUENTIALISM AND THE SIMPLE VIEW

Recall the Simple View, which states that subjective proportionality is reducible to Hurka’s Test – it is a comparison between expected harm and expected good. Some may be concerned that the Simple View is clearly too simple, and a position that only the simplest variants of consequentialism would endorse. Through McMahan’s distinction between narrow and wide proportionality we can already see that objective proportionality is more complex than a simple comparison between harm and good – whether the party to be harmed is responsible or innocent makes a difference – and so surely subjective proportionality must be more complex than a simple comparison between expected harm and expected good.

The distinction between responsible threats and innocent third parties is only one of several complicating factors to which philosophers have argued that proportionality ought to be sensitive. McMahan argues that whether or not harm is intended or foreseen is another major factor that can affect proportionality calculations. David Rodin goes further: he identifies two forms of proportionality (based on liability and lesser-evil justifications, or narrow and wide proportionality in McMahan’s terms), each of which are sensitive to fourteen distinct morally relevant factors.¹⁶ In addition to these concerns, in cases in which we may harm or save more than one person, proportionality may also be sensitive to distributive concerns, such as whether harms are distributed equally or unequally, and whether we can simply aggregate small goods for the many in order to justify large harms to the few.¹⁷ The Simple View, with

¹⁵ See, for example: Jeff McMahan, ‘Proportionality and Necessity in *Jus in Bello*’.

¹⁶ Jeff McMahan, *Killing in War*, 19-29; Rodin, ‘Justifying Harm’.

¹⁷ In cases in which we must choose whom to save, many philosophers believe that no number of small harms, like minor headaches, could add up to outweigh a single life. A related position could be taken regarding harming-to-save cases: you should never kill to prevent any number of minor headaches. On

its comparison between expected harm and expected good, may seem deaf to these concerns.

We can, however, slightly complicate Hurka's Test, and thus the Simple View, in order to take account of these, broadly speaking, non-consequentialist concerns. The best way to do this is to morally-weight the harms in question. So, for example, all else equal, harming an innocent person creates a greater amount of *morally-weighted harm* than harming a culpable aggressor, and so it is harder to justify, requiring more good to be achieved in order to be proportionate. This is what Hurka himself does.¹⁸ According to a suitably amended account of proportionality, objective proportionality asks us to compare morally-weighted harm with good, while the Simple View of subjective proportionality would ask us to compare expected morally-weighted harm (or, perhaps, morally-weighted expected harm) with expected good.

Even on this more sophisticated reading, however, the Simple View contains a core assumption, which, I will argue, renders it irredeemably flawed. As a reminder, that assumption is as follows:

The Separation Assumption: subjective proportionality assessments must consider all of the potential harms that may occur, and then, separately, all of the potential goods that may result. The collective of potential harms must then be compared with the collective of potential goods.

The Simple View, and the related Separation Assumption, are certainly *prima facie* attractive. It is natural to assume that since we are concerned with morally-weighted harm and good at the objective level, we will be concerned with expected morally-weighted harm and expected good at the subjective level. It seems plausible to suppose that we will be concerned by the same things at the objective and subjective

the saving cases see, for example: Alex Voorhoeve, 'How Should We Aggregate Competing Claims?', *Ethics* 125: 64-87; Larry S. Temkin, *Rethinking the Good* (Oxford: Oxford University Press, 2012). I am doubtful of this sort of view. See: Patrick Tomlin, 'On Limited Aggregation,' *Philosophy & Public Affairs* 45 (2017): 232-260.

¹⁸ See 'Proportionality in the Morality of War,' 57-66, where Hurka discusses how to weigh harms to different parties in a war. This approach is recommended for necessity in Seth Lazar, 'Necessity in Self-Defense and War'.

level, and thus that all we need to do to find the right account of subjective proportionality is to modify our account of objective proportionality to include a concern for probabilities.

Nevertheless, I will argue that the Simple View and the Separation Assumption are flawed. The Separation Assumption does not allow us to take the different pairings of harm and good that may come into existence together into account, which in turn rules out considering the objective proportionality or disproportionality of these pairings. Against this, I will argue that there are two different ways in which subjective proportionality ought to take pairings of harm and good into account. First, in Sections IV and V, I will argue that in trying to take account of the moral significance of intended harm, we must take the pairing of the intended harm and the good that it will lead to into account – it matters whether intended harms will lead to objectively proportionate outcomes. Second, In Sections VI and VII, I will argue that we must take account of whether, how seriously, and how likely, our action is to be objectively disproportionate. We can only do so if we reject this Separation Assumption.

IV INTENTIONS UNDER UNCERTAINTY

The first line of objection to the Simple View, and the Separation Assumption which underpins it, concerns intended harm.¹⁹ McMahan and Rodin both make a distinction between intentional and foreseen defensive harms within their accounts of proportionality.²⁰ In doing so, they endorse a weak version of the Doctrine of Double Effect (DDE). On some versions of the DDE, all intentional harming is automatically

¹⁹ It may be thought that since I refer here to intended harm, I have automatically made my claims irrelevant to war, since in war (1) only combatants may be intentionally harmed; (2) all enemy combatants are liable to be killed; (3) combatant deaths are irrelevant to proportionality. Traditional just war theory accepts (3), but I think it is mistaken. We should not kill thousands of soldiers to protect a small strip of unimportant territory (see Jeff McMahan, 'Liability, Proportionality, and the Number of Aggressors,' in Saba Bazargan and Samuel Rickless, eds., *The Ethics of War* (New York: Oxford University Press, 2016)). In addition, I am skeptical of both (1) and (2) as part of the 'deep morality' of war.

²⁰ Some versions of the DDE focus not on intended harm, but on intended effects on a person that foreseeably lead to harm. I aim to be neutral between these formulations here, and references to 'intended harm' can be interpreted as 'intended effects which foreseeably harm'. See: Warren Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Double Effect,' *Philosophy & Public Affairs* 18 (1989): 334-351; Victor Tadros, 'Wrongful Intentions without Closeness,' *Philosophy & Public Affairs* 43 (2015): 52-74.

ruled out as impermissible. On weaker, non-absolutist versions, intentional harm is harder to justify in comparison with foreseen harm of the same magnitude.²¹ Weak versions of the DDE, therefore, essentially demand different moral weightings for intended and foreseen harm. The Weak DDE leads to this claim about proportionality:

Weak DDE: Whether harm is intentional or foreseen affects the moral weighting of harm. There are some levels of good for which it would be proportionate to inflict x units of harm on a person as a side-effect, but would be disproportionate to inflict x units intentionally.

The DDE is often discussed using cases in which we know how much harm will be done, for example the well-known trolley cases. As we shall see here, however, how to incorporate the moral significance of intentions into situations which involve Harm Uncertainty is a complex question.²² Even those who disagree with my conclusions here will, I hope, at least see that it is far from obvious how we should incorporate intentions into subjective proportionality, and a range of possible positions will be developed and explored.

In what follows, I accept the Weak DDE. The Weak DDE will be compatible with the Simple View if we can amend the ‘expected harm’ element of Hurka’s Test to take account of the difference between intended and merely foreseen harms. In the remainder of this section, I will try out a few ways of doing this. Each is found wanting. This leads me to develop an alternative way of capturing the significance of intentional harm within a theory of subjective proportionality – one that is not compatible with the Simple View.

Let us begin by examining the how the Weak DDE works in a case, like those typically presented, in which the defensive agent knows how much harm she will do. Imagine

²¹ For an overview, see: Alison McIntyre, ‘Doctrine of Double Effect’ in *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/double-effect> (2014).

²² It is, furthermore, independent of the already complex question of how to identify ‘intended’ and ‘foreseen’ harms under conditions of uncertainty. I put aside that difficult question here, and focus on it in ‘Accidentally Killing on Purpose’ (unpublished m/s). Here I consider only cases in which there is a clear intended harm (or a clear intended effect on the person which, if it eventuates, will lead to a particular level of harm).

you attack me and I shoot at you, intending and knowing that this shot will kill you. Two things seem relevant here, on the harm side – the amount of harm (death), and that the harm was intended. These are listed as two of Rodin’s fourteen morally relevant factors.²³ Rodin’s approach implies that amount of harm and the intention to harm are two independent factors. A plausible interpretation of this is that when a defensive agent intends harm, this adds a consistent multiplier to the harm done. That would give us this view in certainty cases:

Multiplier: whenever some harm is done intentionally, a set multiplier is applied to the level of harm. Morally-weighted harm = harm x i

Let’s say that $i=5$. For 10 units of intentional harm, the morally-weighted harm would be 50. *Multiplier* may seem plausible in cases where we are certain as to how much harm we will do, and I think this is how most people see the Weak DDE working. The simplest way to convert this view into a view about subjective proportionality, and the one that seems most faithful to Hurka’s Test, is to simply replace harm with expected harm.²⁴ This gives us:

Multiplier (Subjective): whenever some harm is done intentionally, a set multiplier is applied to the level of expected harm. Morally-Weighted Expected Harm = Expected Harm x i

Multiplier (Subjective) is compatible with the Simple View. This view is problematic, however, because a uniform weighting is applied, regardless of how much harm is intended, to all harmful outcomes, whether those are intended or not. Consequently, *Multiplier (Subjective)* cannot differentiate between these two cases:

Finger or Intentional Death: Attacker threatens Victim with a broken finger. Victim responds, intending to kill Attacker by shooting him, in order to avert her finger being broken, but with only a small chance of success. If Victim misses, the loud noise will frighten and thus distract Attacker such that he will

²³ Rodin, ‘Justifying Harm’, 81, 92-95.

²⁴ To be clear, while I used Rodin’s work to generate *Multiplier* as a plausible operationalisation of the Weak DDE for certainty cases, I do not mean to suggest that *Multiplier (Subjective)* is his considered view under uncertainty.

break Victim's little finger rather than her index finger. This will be a slight harm to the Attacker, and result in a slightly lesser harm to the Victim.

Finger or Accidental Death: Attacker threatens Victim with a broken finger. Victim responds, intending to shoot into the air in order to frighten Attacker, thus distracting him so that he will break Victim's little finger rather than her index finger. She is almost certain to succeed in doing this. However, Victim knows that there is a very small chance that she will accidentally kill Attacker. This will prevent the attack altogether.

If the probabilities of killing and scaring the attacker in the two cases are identical, the expected harm and the expected good will be identical. And since harm is intended in both cases, then morally-weighted expected harm would also be identical. Therefore, Multiplier (Subjective) would treat these two cases identically.

Yet these two cases seem importantly different, due to the differing intentions. At the very least, I think there will be some levels of expected harm and good at which it is subjectively proportionate to shoot in *Finger or Accidental Death*, but not *Finger or Intentional Death*. More strongly, I am inclined to say that Victim acts impermissibly in *Finger or Intentional Death* whatever the probabilities: it is potentially permissible to risk death, but not to try to cause it, in this scenario, when a broken finger is all that is at stake.

An alternative way to try to incorporate the Weak DDE within the Simple View is to limit the additional weighting for intended harm only to harm which is intended. Expected morally-weighted harm would be calculated by looking at each potentially harmful outcome, but only adding additional weight to the intended harm. Imagine only Potential Harm 1 (PH₁) is intended. We'd then find the Expected Morally-Weighted Harm in this way:

Limited Weighting (Subjective): Expected Morally-Weighted Harm = (PH₁ x *i* x Probability) + (PH₂ x Probability) + (PH₃ x Probability)... + (PH_n x Probability).

This view can potentially differentiate between our two Finger or Death cases: in Finger or Intentional Death the additional weighting would be applied to death, whilst in Finger or Accidental Death the additional weighting would be applied to scaring the assailant with a loud noise. The problem with Limited Weighting (Subjective) is that the additional weighting for intentional harm is discounted by probability. This is counter-intuitive for several reasons. First, when the chances of succeeding with the intentional harm are very low (as in Finger or Intentional Death), hardly any weighting will be added. This makes it a very permissive view. Second, our two Finger or Death cases could still be seen as equivalents, since the additional weighting applied to scaring the assailant in Finger or Accidental Death will be applied to a very likely outcome. Indeed, on this view, trying to kill an attacker may produce less expected morally-weighted harm (and thus be easier to justify) than trying to mildly harm him, even when the potential harms, goods, and their probabilities are kept constant.²⁵

Third, this approach doesn't seem to capture the normative significance of the intentions of the defensive agent in the right way. When we intend something, it is given special normative significance because of our special relationship to that outcome. It isn't merely one outcome among the pack. Discounting the additional weighting for probability seems to mix together two independent moral concerns – what am I trying to make happen?; and what is likely to happen? It matters what the defensive agent is aiming for, and the significance of that cannot be fully accounted for in a calculation which discounts that significance according to the likelihood of it occurring.

Limited Weighting (Subjective) discounts the significance of intended harm for probability, whereas Multiplier (Subjective) applies a uniform multiplier to all harms.

²⁵ For example, imagine that death is 90 units, and mild harm is 1 units. Imagine that there is a 1% chance of death, and 99% chance of mild harm, and that intended harm is weighted three times more heavily than foreseen harm. If death is intended, then Limited Weighted (Subjective) would calculate expected morally-weighted harm as $(90 \times 0.01 \times 3) + (1 \times 0.99) = 3.69$. If mild harm is intended, then the view would calculate expected morally-weighted harm as $(90 \times 0.01) + (1 \times 0.99 \times 3) = 3.87$. Therefore, when everything else is kept constant, intending a harm that is ninety times more serious can result in less expected morally-weighted harm.

An alternative is to apply additional weight that is sensitive to how much harm is intended:

Intended Harm Multiplier (Subjective): Morally-Weighted Expected Harm = Expected Harm x IH (where IH is a multiplier that increases with level of harm intended).

Note that this view, like Multiplier (Subjective), will increase all potential harms by the same multiplier, regardless of whether they are intended or merely foreseen. This view, however, can differentiate between the two Finger or Death cases. When probabilities of outcomes are left constant, the morally-weighted expected harm of intending death will be greater compared with intending to scare someone. Furthermore, this view does not discount the importance of intended harm by its probability, and so Finger or Intentional Death may still be very hard to justify (though not impossible). Since it has these entailments, this view seems preferable to both Limited Weighting (Subjective) and Multiplier (Subjective), and it is compatible with the Simple View.

Some may wish to endorse this view. But, if this is the right view, this tells us something important about the moral significance of intentions – it isn't that magnitude of harm matters, and, separately, whether or not the harm was intended matters, as Rodin's separation of these two concerns may seem to suggest. Rather, the magnitude of intended harm matters. The distinction between these two interpretations of the significance of intentional harm is lost when we focus on cases in which you are certain as to how much harm you will do.

However, Intended Harm Multiplier (Subjective) would be very demanding in cases with large numbers. We are often tempted to think that the Weak DDE gives us a consistent ratio of harm to good (all else equal), so we think that we can read from, for example, a trolley case to an act of war: if it is permissible to intentionally kill one person in order to achieve some level of good, we think it will be permissible to intentionally kill 100 people in order to achieve some level of good 100 times greater. Intended Harm Multiplier (Subjective) denies this. It says that killing 100 people

would not necessarily be justified. Since the intended level of harm (100 deaths) is so much greater, IH would be very large, and so the morally-weighted harm would be much greater than 100 times that of intentionally killing one.²⁶

All of the views discussed thus far are compatible with a suitably amended Simple View, and the Separation Assumption. Whichever we go for, however, we are left with two problems. The first is that they all allow that Finger or Intentional Death may be permissible, if the chances of succeeding fall low enough. It could be subjectively proportionate to try to kill someone in order to avoid a broken finger. I will return to this below, but it is a conclusion that many will want to resist.

Second, none of the views articulated thus far can differentiate between the following two cases:

Intentional Death for a Fingernail: Attackers 1 and 2 threaten Victim with a broken fingernail. If Observer does nothing, the Attackers will break the fingernail, and also kill Bystander as a side-effect. Observer responds by trying to shoot Attacker 1 in the chest, in order to kill him, thus preventing the attack on Victim (but allowing Bystander to die). She foresees that if she misses, she will kill Attacker 2, thus saving Bystander's life (but Victim's fingernail will be broken). There is 50% chance that she will succeed in killing Attacker 1.

Intentional Death for a Life: As *Intentional Death for a Fingernail*, except this time Observer tries to shoot Attacker 2 in the chest, in order to kill him, thus saving Bystander's life (but allowing Victim's fingernail to be broken). She foresees that if she misses, she will kill Attacker 1, thus saving Victim's fingernail (but Bystander will die). There is 50% chance that she will succeed in killing Attacker 2.

In these cases the expected harm, the level of harm intended, and the probability of

²⁶ Imagine IH were 3 times the level of intended harm, and one death = 1 unit of harm. In that case, the morally-weighted harm of intentionally killing one person is 3 units ($1 \times (3 \times 1)$). The morally-weighted harm of intentionally killing 100 people is 30,000 units ($100 \times (100 \times 3)$).

the intended harm eventuating are identical. Therefore, each of the views canvassed above would give us the same level of expected morally-weighted harm. And the expected good is also equal in both cases. All of the above ways of incorporating the significance of intended harm into the subjective proportionality calculation would see these two cases as equivalent.²⁷ Intuitively, however, there is a moral distinction between these two cases, one that renders Intentional Death for a Fingernail, at the least, much harder to justify.

However, in order for these cases to show that the views canvassed above about subjective proportionality are incorrect, I need to show that the moral distinction between Intentional Death for a Fingernail and Intentional Death for a Life is one to which subjective proportionality in particular ought to be sensitive.

It is easy enough to show that the difference between the cases involves proportionality. The only difference between the cases is the way that the intended harm matches up with the greater or the lesser good. In Intentional Death for a Fingernail, the intended death will do very little good. Precisely what seems to matter about this case is that the intentional harm will be objectively disproportionate: while there is uncertainty as to whether the intentional harm will come about, there is certainty that the intentional harm, if it eventuates, will be objectively disproportionate. In Intentional Death for a Life the intended harm is objectively proportionate, though an objectively disproportionate outcome is risked. So, the difference between the cases concerns the objective (dis)proportionality of the intended harm. Nevertheless, some may be tempted to say that while that is indeed the difference between the cases, it doesn't follow that this shows that the moral error committed in Intentional Death for a Fingernail is one of subjective proportionality.

²⁷ The probabilities of the two potential pairings of harm and good need only both be 50% in order for Limited Weighting (Subjective) to fail to differentiate between the two cases. Both Multiplier (Subjective) and Intended Harm Multiplier (Subjective) would fail to differentiate between the cases whatever the probabilities of the two pairings within the cases, provided that the probabilities of the pairings were identical across the two cases (for example, a 30% chance of killing-for-a-fingernail, and a 70% chance of killing-for-a-life).

There are two claims that might be made here. One is that the necessity principle is the principle that is implicated here. In *Intentional Death for a Fingernail*, since Observer aims at an objectively disproportionate outcome, that should add, we might think, to the morally-weighted expected harm. Since Observer could have aimed to save a life, thereby giving herself a proportionate aim, Observer does more morally-weighted expected harm than she needs to when she gives herself a disproportionate aim. As such, her conduct is condemned by the necessity principle. But if morally-weighted expected harm is increased, this ought to matter for subjective proportionality as well, even if it also matters for necessity, so the response seems, ultimately, to concede that both subjective proportionality and necessity calculations will be affected by the intended harm being objectively disproportionate.

Alternatively, it might be claimed that the fact that Observer's intended harm will be objectively disproportionate in one case, but not the other, shows that the problem is not with subjective proportionality, but rather something like a 'right intention' principle. But, if so, this is revealing about the right intention principle. The right intention principle is usually thought to require that wars are fought 'for the right reasons' – that is, for a just cause. But protecting someone's fingernail is a just cause – it would justify very mild levels of harm. The problem in *Intentional Death for a Fingernail* is that the amount of harm is excessive. So, even if this response is granted, it is only because the right intention principle turns out to be a proportionality principle. The question we ask is not 'are you doing this for acceptable reasons?' but 'are your ultimate aims proportionate to the harms you intend?' Whatever we call this, this is a question of proportionality that arises under conditions of uncertainty.

If I am right about all of this, it matters for subjective proportionality whether or not the intended harm leads to an objectively proportionate outcome.²⁸ However, we

²⁸ I have only considered cases here in which we can identify some particular level of intended harm, or in which the intended effect on the person that will lead to a particular level of harm (see n. 24 above). Sometimes, however, we intend to perform a harmful action which could lead to a variety of harmful outcomes. Imagine, for example, that I intentionally throw you down the stairs, averting your attack, but I do so knowing that you will either break your arm or your leg. According to my account here, the intended harm must be assessed for its proportionality. But in this kind of case, what figure should we use for the intended harm? This is a very important but difficult question, and what position we take on

cannot ignore the other, unintended, outcomes. Imagine that in Finger or Accidental Death shooting in order to scare was overwhelmingly likely to unintentionally kill. Whilst this act would pose no problem from the point of view of intentions, it would certainly be subjectively disproportionate and should not be performed, due to the likely but unintended effects. There are therefore (at least) two factors that an approach to subjective proportionality must take into account – the objective proportionality of the intended harm; and the full range of potential outcomes (including their likelihoods of occurring).

These two factors could be combined in two different ways. The first is that they could be independently necessary (and, perhaps, jointly sufficient) conditions for a harmful act being subjectively proportionate. That is, acts of intentional harm must pass two tests – one asks whether the intentional harm would lead to an objectively proportionate outcome; while a second takes account of the full range of possible outcomes, and their probabilities. The second way to combine these concerns is to make intentional harms that lead to objectively disproportionate outcomes add to the morally-weighted expected harm of an action, which will be measured against the expected good. This gives us these two views:

Independent Justification: Intended harms are subject to two independent tests. First, does the intended harm lead to objectively proportionate outcomes? Second, is the overall level of risk acceptable, given all the outcomes that could eventuate, and their probabilities of eventuating?²⁹

Incorporation: Intended harms that lead to (or risk leading to) objectively disproportionate outcomes are harder to justify. That is, they add to the

it will most likely be informed by wider commitments concerning how to understand intentions and their moral significance. These are debates and issues I cannot tackle here, for space reasons. Therefore, I will briefly outline three stances we might take on this question. First, we could take the intended harm to be the expected harm of the intended effect. Second, we could take the intended harm to be the minimum possible harm that will achieve the intended good. Third, we could assess the intended effect's Expected Objective Disproportionality (see Section VI for an explanation of this measure).

²⁹ This second test could, on the argument thus far, be Hurka's Test. I will show in Sections VI and VII why we should reject that role for Hurka's Test.

morally-weighted expected harm of an action. This morally-weighted expected harm must be compared with the expected good of acting.

Endorsing either of these views involves rejecting the Simple View and the Separation Assumption. Independent Justification requires two independent proportionality calculations for subjective proportionality, so subjective proportionality cannot simply be about comparing morally-weighted expected harm with expected good. And on both accounts, it matters whether the intended harm leads to (or may lead to) an objectively disproportionate outcome. Since objective proportionality concerns the pairings of harm and good that eventuate together, any such concern is incompatible with the Separation Assumption.

I can see the attractions of both views, but I am inclined toward Independent Justification. To see why, imagine in *Finger* or *Intentional Death* that Victim's attempt is successful, and she kills Attacker. Now consider her trying, *ex post*, to justify her actions, by appeal to her *ex ante* situation. She recognizes that the harm imposed turned out to be objectively disproportionate. But, she says, the morally-weighted expected harm was far lower – a very low risk of death was imposed upon attacker. In my view, this cannot justify her action. This was not a roulette wheel – her action did not simply happen to turn out to be objectively disproportionate. Rather, she tried to kill and succeeded, knowing that that would produce an objectively disproportionate outcome. The fact that she might have been unsuccessful cannot justify the harm she has caused, any more than a successful murderer can point toward the fact that he was extremely unlikely to succeed, and that other potential outcomes might have been beneficial to his victim.

I prefer Independent Justification because it rules out such appeals to less harmful but unintended outcomes when the intended harm is objectively disproportionate. The Incorporation view, on the other hand, merely puts a thumb on the scales for intended objectively disproportionate harms, making them harder to justify, but it does not rule them out. Thus, objectively disproportionate intended harms can still, on that view, be

made subjectively proportionate by unintended outcomes: one is permitted to try to bring about a harm that one would not be permitted to bring about if one knew one would be successful.

Out of the range of possible harms that could stem from Victim's action, she has a special moral relationship to a particular one: the intended harm. That intended harm must figure in her justification of her action. And in order to justify that harm, she must justify it in terms of the outcome(s) it will lead to, and in particular the objective proportionality of those outcome(s). If she is not trying to bring about an objectively proportionate outcome, her action is not subjectively proportionate. Only Independent Justification takes this stance.

V JUSTIFYING INTENDED HARM

I have argued that, when harm is intended, it matters for subjective proportionality whether the intended level of harm leads to objectively disproportionate outcomes. Thus far we have considered cases in which we are uncertain as to whether the intended harm will eventuate, but certain as to what would follow from that harm eventuating. These cases, therefore, avoided some difficult questions about exactly what we should care about when we care about whether or not the intended harm 'leads to' objectively disproportionate outcomes: that the intended harm might lead to an objectively disproportionate outcome; that it carries an unacceptable risk of an objectively disproportionate outcome; or that an objectively disproportionate outcome is intended? In other words, I have claimed that intended harm must be subject to a proportionality calculation. But we must have some measure of goodness to put on the other side of that calculation.

To help make things clear, consider the following the case:

Certain Death, Uncertain Defence: Attacker threatens Victim's life. Victim has only one defensive option available to her – to shoot attacker and kill him. There are three possible good outcomes of killing attacker, compared with doing nothing: that Victim will prevent a minor harm to an innocent bystander, Alice, but will not save her own life; that Victim will save her own life; that

Victim will save the lives of Barbara and Charlotte, who are innocent bystanders, but not her own life. Victim hates Alice, Barbara, and Charlotte and would actively prefer them to be harmed. She in no way intends to save them. She intends to save her own life, and only shoots with that intention in mind.

Victim's act, as an intentional harm, must be justified as subjectively proportionate, and in order to do so, I have argued, we must look at whether the intended harm leads to objectively proportionate goods. But which of the goods on offer should we measure the intentional killing against?³⁰

There are two broad approaches we can take here. Either we must, in some way, take account of all of the goods that the intended harm will produce (including both the objectively disproportionate outcome of saving Alice, and the objectively proportionate outcome of saving Barbara and Charlotte)³¹, or we should measure the intended harm only against the good outcome that the defensive agent intends it to produce.

My own view, which I will tentatively defend here, is that intended harms should be defended as proportionate by comparing them with the intended good (namely, in this case, saving Victim's life only). The main question we have is whether goods which may be the product of the intended harm but are not the intended goods can help to justify the intended harm qua intended harm. (Recall that another calculation, or another element of the subjective proportionality calculation, will take account of the full range of possible outcomes).

I think that unintended goods cannot help to justify intended harms. When someone is looking to harm others, we must essentially ask the putative harmer (or demand that she asks herself) two questions: First, what are you trying to do? And second,

³⁰ For related discussion, see Rodin, 'Justifying Harm', 94-95.

³¹ This could be achieved by measuring the expected good of the intended harm, or the Expected Objective Disproportionality of the intended harm. I introduce the notion of Expected Objective Disproportionality in the next section.

what do you think will happen? In order to justify herself, both answers need to be proportionate. When harm is intended, in answering the question ‘what are you trying to do?’ the answer must involve a particular pairing of harm and good – if she is to be justified, then she must be trying to bring about harm in order to bring about some good outcome. Therefore, she should justify that pairing, which means (among other things) that it must be objectively proportionate. When an intentional harm may produce unintended goods, I think we should say the same thing as we say about unintended harms – you can’t appeal to the fact that you might fail in what you are trying to achieve in order to justify trying to bring about a particular outcome.

While I have tried, here, to briefly defend the idea that intended harms must be justified as proportionate to intended goods, I am more confident of the more general claim that it matters, for subjective proportionality, whether or not intended harms lead to objectively disproportionate outcomes. Therefore, those who reject the argument of the present section, should not necessarily reject the more general argument that the objective (dis)proportionality of intended harms has moral significance relevant to subjective proportionality, and thus that the Simple View and Separation Assumption must be rejected.

VI EXPECTED OBJECTIVE DISPROPORTIONALITY

The Simple View states that subjective proportionality is reducible to Hurka’s Test. Thus far, I have argued that this comparison is not sufficient for subjective proportionality when the harm is intended, since it matters whether or not intended harms lead to objectively proportionate outcomes. However, along the way we met various other ways of incorporating the significance of intentions into subjective proportionality which were compatible with the Simple View. Some readers may have been tempted to accept one of those, and thus to affirm the Simple View and the Separation Assumption. Others might deny the moral significance of intentions, and thus deny that what I have said thus far has any relevance whatsoever to subjective proportionality.

Even those who accept my analysis in full thus far may see space left for the Simple

View, or Hurka's Test at least. First, they may think the Simple View remains the right view when intended harms are not involved. Second, the argument thus far has led us to Independent Justification: a position which requires a proportionality test for intended harms and a second test that takes account of the full range of possible harms and goods. For all we have said thus far, Hurka's Test may be the right account of this second test.

I think, however, that there is a deeper problem with the Simple View and the way that Hurka's Test handles the range of potential harms and goods. Therefore, I think the Simple View fails even if my arguments around intentional harms are rejected, and that Hurka's Test should play no role in subjective proportionality. Consider a case of non-intentional harming in which each defensive option, if performed, is certain to prevent an attack completely. However, each comes with different risks concerning the morally-weighted harm that may occur. Let's imagine that -10 units of morally-weighted harm will be objectively proportionate. The defensive agent has the following defensive options.

Option 1 – Certainty: -8 units of morally-weighted harm for sure.

Option 2 – Uncertainty 1: 50% chance of -10 units of morally-weighted harm; 50% chance of -6 units of morally-weighted harm.

Option 3 – Uncertainty 2: 50% chance of -14 units of morally-weighted harm; 50% chance of -2 units of morally-weighted harm.

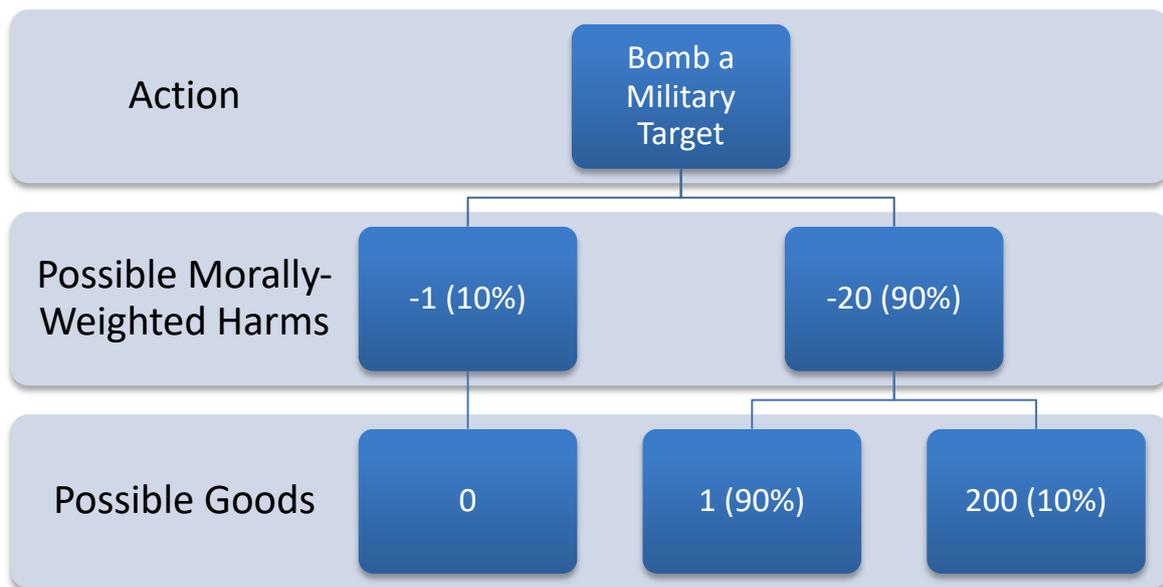
The Simple View would view each of these options as morally identical. Each will do (for sure) the same amount of good, and each has an expected morally-weighted harm of -8 units. However, I think there is a morally important distinction between Options 1 and 2 on the one hand, and Option 3 on the other. In Option 1 the harmful outcome is certain, and is within the objectively proportionate limit – it is clearly subjectively proportionate. In Option 2 there is uncertainty as to how much harm will be done, but since each potential harmful outcome will be objectively proportionate, it must be subjectively proportionate. In Option 3, however, there is a serious risk that the harmful outcome will be objectively disproportionate. The amount of good that will be done can only justify up to -10 units of harm, and yet it is as likely as not that -14 units

of harm will be done. This outcome is one that we would be forbidden to bring about knowingly. If this greater harm comes about, people will be harmed in a way that wrongs them – the good outcome is not good enough to justify the harms inflicted upon them.

This case is designed to illustrate that we are concerned about proportionality under uncertainty because we are concerned that our harmful action may turn out to be objectively disproportionate.³² If all of our options are objectively proportionate, the fact that we are uncertain about which will eventuate is not a concern from the point of view of proportionality. The Simple View, relying on its expected harm-expected good comparison, cannot capture our worries about objective disproportionality because objective disproportionality comes about when a particular combination of harm and good come into existence together, and the Simple View takes no account of these combinations.

In order to put issues of intended harm to one side, let us focus on a case in which all of the potential harms are merely foreseen. Consider this outcome tree:

³² The only account I know which takes the risk of objective disproportionality into account is Adil Ahmad Haque's. However, Haque's account, in my view, suffers from some crucial defects. He argues that what matters is whether the act will 'probably' (i.e., will more likely than not) result in an objectively proportionate outcome. I disagree with this test for two reasons. First, it fails to take into account *how* disproportionate our action may be: all objectively disproportionate potential outcomes are simply treated alike. Second, any lowering of the chances of a disproportionate outcome are ignored once we get over the 'more likely than not' threshold: a 51% chance of an objectively proportionate outcome is just as good as a 100% chance. Haque couples this principle with test similar to Hurka's. See Adil Ahmed Haque, *Law and Morality at War* (Oxford: Oxford University Press, 2017), 195-199.



In this case, the expected morally-weighted harm is -18.1 ³³, and the expected good is 18.81 .³⁴ According to the Simple View, this would be subjectively proportionate (assuming, for simplicity, a $-1:1$ proportionality ratio³⁵). Expected harm and expected good are calculated by looking ‘horizontally’ – first at all the possible harms, and then all the possible goods. But if, as I have argued, we should be worried about the particular combinations of harm and good that may eventuate, we need to look at the ‘vertical’ combinations. Here we see that the most likely combination (-20 units of harm for 1 unit of good, which is 81% likely to occur³⁶) is seriously disproportionate. And there is a 91% chance that the act will be objectively disproportionate. Almost all of the ‘expected good’ is coming from a 9% chance of a large amount of good. In other words, the beyond proportionate goods of that (unlikely) outcome are allowed to fully compensate for the proportionality shortfalls of more likely options: there is nothing in Hurka’s Test that even attaches any additional weight or concern to objectively

³³ Expected Harm = $(-1 \times 0.1) + (-20 \times 0.9) = -18.1$.

³⁴ The good outcomes are 0 , 1 , and 200 units of good. 0 units is 10% likely to occur. 1 unit is 81% likely to occur (since there is a 90% chance of -20 units of harm occurring, and a 90% chance that 1 unit of good will follow from that). 200 units is 9% likely to occur (since there is a 10% chance of it following from the 90% chance of -20 units of harm). Therefore, Expected Good = $(0 \times 0.1) + (1 \times 0.81) + (200 \times 0.09) = 18.81$.

³⁵ Recall that the harm is morally-weighted, and so it certainly does not follow that a $-1:1$ ratio here implies that one can proportionately cause one unit of harm in order to avoid one unit of harm.

³⁶ Since 1 unit of good has a 90% chance of being produced if -20 units of harm are produced, of which there is a 90% chance.

disproportionate outcomes. This does not seem to take the idea of proportionality as a limit seriously – it is supposed to be a ceiling on the harm we are allowed to impose for a given level of good, not an average we should not exceed.

We can take ‘vertical’ combinations into account if we look at each combination in turn, and assess its objective disproportionality. In order to assess a combination’s objective disproportionality, we need to look at two things. First, is the harm objectively proportionate or disproportionate? Second, if it is objectively disproportionate, how seriously objectively disproportionate is it?

How do we measure how seriously objectively disproportionate an outcome is? There are, broadly-speaking, two ways in which we can do so. One is that we say that any death or injury that is part of an objectively disproportionate outcome is objectively unjustified, and that all objectively unjustified deaths and injuries count against an outcome. So, if -5 units of harm would be an objectively proportionate outcome, and an outcome involves -6 units of harm, we should count all -6 units of harm as objectively unjustified, and this is the measure of how seriously objectively disproportionate the outcome is. An alternative is to measure *surplus harm*. So, on the assumption that -5 units of harm would have been objectively proportionate, in an outcome where there is -6 units of harm, there is -1 unit of surplus harm – that is how far past the objectively proportionate maximum the harm is.

Which is the right approach will depend at least in part on how we intend to use the resulting measure, an issue to which I will return shortly. I find the surplus harm version more plausible, since it allows that there is a difference between causing 100 deaths when the objectively proportionate limit is 99, compared with causing 100 deaths when the objectively proportionate limit is 5. The ‘unjustified harm’ approach can’t recognize a difference between these two cases, since in both there are 100 objectively unjustified deaths.

For now, then, I will use the surplus harm approach. To see how this works, let us look at the above outcome tree. The first combination (-1 harm, 0 good) is -1 unit of harm

over the objective proportionate maximum. The second (-20, 1) is -19 units of harm over the objectively proportionate maximum. The third (-20, 200) is objectively proportionate. In fact, it is a long way under the objectively proportionate maximum harm (which is -200 units), but from an objective proportionality perspective this is irrelevant. Recall that proportionality is a limit – there are not extra ‘proportionality points’ for doing less harm than the maximum allowed. Since the outcome is objectively proportionate, it should score a 0 – there is no objective disproportionality. We can then assess the ‘Expected Objective Disproportionality’ of the action as a whole by using these scores and the probability of each combination. In this case, the overall Expected Objective Disproportionality of acting would be 15.49.³⁷

I am convinced that this calculation captures something that is morally relevant in determining what we are permitted to do. It assesses the expected magnitude of how seriously objectively disproportionate a harmful action is likely to be, and the lower the number the better. This leads to some attractive normative entailments that differ from those of Hurka’s Test. Depending on how it is used (on which, see below) Expected Objective Disproportionality would either rule out, or make harder to justify, risky wars – wars that are almost certain to be objectively disproportionate, but carry some very low chance of being better-than-proportionate. Consider a war that only just passes Hurka’s Test, but is almost certain to be seriously objectively disproportionate. Allowing Expected Objective Disproportionality to play a role would allow us to rule out such a war, whereas the Simple View has no resources to do so.

VII USING EXPECTED OBJECTIVE DISPROPORTIONALITY

Whilst the Expected Objective Disproportionality measure ought to be used in some way, I have not, as yet, said how. I will lay out some options here, and explain why I favour the view that I do.

³⁷ The first outcome (-1,0) is 10% likely to occur, and has a OD score of 1, therefore its Expected Objective Disproportionality (EOD) = $(1 \times 0.1) = 0.1$. The second outcome (-20,1) has an OD score of 19, and is 81% likely to occur. Therefore, its EOD = $(19 \times 0.81) = 15.39$. The final outcome (-20,200) has an OD score of 0, since it is objectively proportionate. Therefore, overall the action’s EOD = $15.39 + 0.1 + 0 = 15.49$.

1. No Expected Objective Disproportionality

Proportionality is a limit. If an action is objectively disproportionate, it ought to be ruled out. If we are to take the idea of proportionality as a limit seriously, we might think that if an action carries even a risk of being objectively disproportionate, it should also be ruled out. Therefore, an Expected Objective Disproportionality of greater than zero would be subjectively disproportionate.

This view, however, is clearly too restrictive. Consider a scenario in which I can kill one person. This is almost certain to end a war, saving many thousands of lives. There is also a tiny chance that it will not do any good. If so, there is some chance the act will produce an objectively disproportionate outcome. Yet surely this act should be allowed, and should be regarded as subjectively proportionate.

2. Minimize Expected Objective Disproportionality

We could instead minimize Expected Objective Disproportionality: pick the option with the lowest score. This seems sensible, since if we are worried about objective disproportionality, we should seek to minimize it.

The central problem with this approach is that it may end up ruling out any (course of) action with any Expected Objective Disproportionality. Doing nothing – allowing oneself (or others) to be harmed – does no harm. It is therefore certainly proportionate, and so has no Expected Objective Disproportionality. Therefore minimizing Expected Objective Disproportionality may require us to do nothing whenever doing something has some risk of Expected Objective Disproportionality, no matter how slight. This is implausible for the same reasons as the previous position.

This might be avoided if we exclude ‘doing nothing’ as an option. However, if we do this, and then use Expected Objective Proportionality to compare our defensive options, picking the one with the least Expected Objective Disproportionality, this has three problems. First, it seems to assume that one of our defensive options will be permissible, or at least subjectively proportionate, when we should assume no such thing. Second, and relatedly, this seems to misunderstand the job of proportionality. It

is proportionality's job to tell us which (if any) of our defensive options are potentially permissible by comparing them with doing nothing. It is necessity's job to compare our different defensive options.

The third problem here is that by looking only at an action's Expected Objective Disproportionality, then more and more expected good could never speak in favour of an action with greater Expected Objective Disproportionality. Consider a case in which there are several (only just) objectively disproportionate outcomes, and one objectively proportionate outcome. It surely makes a difference whether the objectively proportionate outcome is one in which one person will die and twenty will be saved, or one in which one person will die and two million will be saved.

Minimizing Expected Objective Disproportionality would fail to take account of this.

3. Morally-Weighted Expected Harm+

Another option is to add an action's Expected Objective Disproportionality to its morally-weighted expected harm on the negative side of the subjective proportionality calculation, and compare these to its expected good. In other words, each potential death counts against an action, but so does the Expected Objective Disproportionality. In Parfit's terms, we must add the non-deontic-badness of the deaths to the deontic badness of the objectively disproportionate outcomes.³⁸ Together, these must be justified as proportionate to the expected good of acting. It is worth remarking that even though this approach resembles the Simple View, it is not compatible with it, as we defined it at the outset. The additional weighting from Expected Objective Disproportionality is generated by particular pairings of harm and good, which violates the Separation Assumption.

One problem with this approach is that it 'double counts' the goods of the objectively disproportionate possible outcomes, at least if we use the 'surplus harm' approach to Expected Objective Disproportionality. These goods count, first, in working out how seriously disproportionate a given harmful outcome is, and then again in helping to

³⁸ Derek Parfit, *On What Matters*, Vol. 3 (Oxford: Oxford University Press, 2017), 351.

offset the Expected Objective Disproportionality. Imagine, for example, that my act has two possible outcomes – it will either save John at objectively proportionate cost, or Elis at objectively disproportionate cost. Saving John would only count toward the expected good of the act. But saving Elis would count twice. First, it would be used to calculate how seriously objectively disproportionate that outcome would be, (so, if saving Elis would justify 100 units of harm, but would actually cause 150 units of harm, that potential outcome would be objectively disproportionate by -50 units). Second, saving Elis (along with saving John) would be used as a potentially good outcome in calculating the expected good of acting. In other words, saving Elis would both make the bad outcome less bad, and be considered, independently, as a good outcome.

Another problem with this approach is that, since it allows that there is some ‘tolerable level’ of Expected Objective Disproportionality, it would allow that an action in which each and every potential outcome was objectively disproportionate could be subjectively proportionate, provided the Expected Objective Disproportionality were low enough. This is clearly unacceptable – if all of our options are objectively disproportionate, we ought not to proceed. This problem could be fixed by adding an additional rule, which is that in order to be subjectively proportionate, an action must at least have a chance of being objectively proportionate.

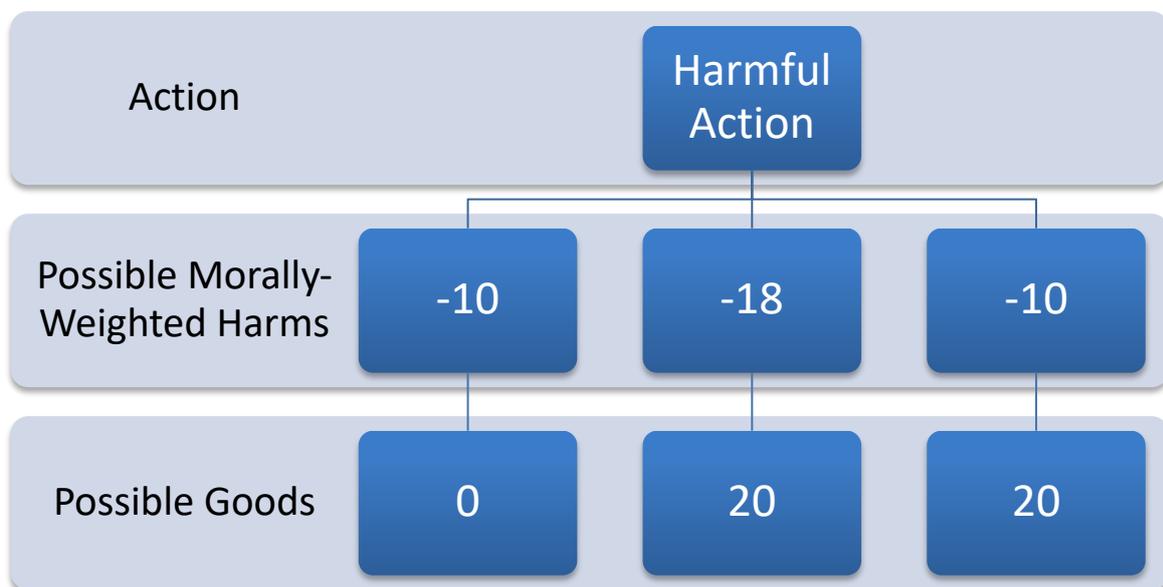
4. Proportionate Disproportionality

An alternative way to take Expected Objective Disproportionality into account is to take Expected Objective Disproportionality to be something that stands in need of justification. This seems to demand the development of a ‘meta-proportionality’ – how much objective disproportionality can be risked? In order to know this, we need to know what Expected Objective Disproportionality ought to be compared with.

Proportionality, in the simplest terms, is about comparing bad things with good things. Here, the bad things are the objectively disproportionate potential outcomes. Just how bad they are, collectively, is measured by Expected Objective Disproportionality. 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. Here we have two choices. We could compare Expected Objective Disproportionality with the Expected Goodness of Objectively Proportionate Outcomes, or with the Expected Surplus Goodness of Objectively Proportionate Outcomes.

To see the difference between these two measures, consider this ‘outcome tree’:



Let us assume that -1:1 is the proportionality ratio. The first outcome (-10:0) is objectively disproportionate. Risking this outcome must be justified. The second outcome (-18:20) is objectively proportionate. The third outcome (-10:20) is also objectively proportionate. Our question, then, is whether the possibility of the second and third outcomes can justify risking the first. There are two ways to think about this. First, we could look at the goods that will be produced in the objectively proportionate outcomes (the Expected Goodness of the Objectively Proportionate Outcomes), in which case, both outcomes are the same, producing 20 units of good each. Or, second, we could look at the surplus good – that is, the goods beyond the objectively proportionate minimum (the Expected Surplus Goodness of the Objectively Proportionate Outcomes), in which case the third outcome looks better than the

second.

I think that which of these is the right view depends on whether the putative justification for causing or risking harm is a liability justification (narrow proportionality) or a lesser-evil justification (wide proportionality). That is, when the person has made themselves liable to harm, so long as the goods are bought at the cost of proportionate harm, the level of harm does not count against a good outcome, so we should look at Expected Goodness of the Objectively Proportionate Outcomes. Therefore, the second and third outcomes would be judged identically: both produce 20 units of goodness at ‘acceptable cost’ – the harms in question do not run up against rights. However, when we are harming innocent persons, even acceptable costs are regrettable in a way that they are not when done to the liable. Therefore, these harms still count against an outcome, and so we should look at the Expected Surplus Goodness.

To put things more concretely: imagine it is objectively proportionate for me to cause harms up to paralyzing an attacker if this will avert the threat that they pose. Outcomes in which I avert the threat at the cost of a broken finger and at the cost of a broken leg should be viewed identically for the purposes of seeing whether objectively disproportionate outcomes (e.g., those in which the attacker dies) are worth risking. But if I am harming an innocent person to avert some threat, it very much matters whether their leg or finger is broken in order to avert the threat, even if both are objectively proportionate.

Expected Objective Disproportionality ought to be taken into account in subjective proportionality calculations. Of the ways in which it might be taken into account canvassed above, Morally-Weighted Expected Harm+ and Proportionate Disproportionality seem to be the most plausible contenders. Whichever we choose, however, in taking account of pairings of harm and good, we have gone beyond the Simple View and the Separation Assumption.

My own view is that Proportionate Disproportionality best captures the way in which

objective proportionality is a limit, without leading to the implausible entailments of No Objective Disproportionality. Proportionate Disproportionality says that 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 proportionate potential outcomes. Although it does not view potential objective disproportionality as ruling out an action, it does take the line between objective proportionality and objective disproportionality to be a morally significant threshold. This also means that an action which could only produce an objectively disproportionate outcome could never be viewed as subjectively proportionate.

Adopting this calculation is compatible with still leaving some role for Hurka's Test in our account of subjective proportionality. However, once we have this measure, it is very hard to see what additional work the direct comparison between expected harm and expected good can do. For this reason, I reject Hurka's Test altogether: subjective proportionality isn't about directly comparing expected harm with expected good at all. Instead, we must look at pairings of harm and good, paying special attention to some of them.

VIII CONCLUSIONS

In this paper, I have explored the important topic of subjective proportionality. If we want to be able to perform proportionality calculations in the real world, we need to know how to do them under conditions of uncertainty.

I have tried to show the following. First, it is not obvious how we should fold the moral significance of intentions into subjective proportionality – several possible views were developed and discussed. Second, the most plausible views take into account whether the intended harm leads to objectively proportionate or disproportionate outcomes. Third, a particularly attractive variant of that stance is that it is an independent subjective proportionality criterion that the coupling of harm and good intended is objectively proportionate. Fourth, the risk of objective disproportionality should be taken into account: we should take account of Expected Objective Disproportionality. Fifth, this can be done in several ways, but one especially attractive

way is for Expected Objective Disproportionality to be measured against the Expected Goodness of Objectively Proportionate Outcomes. How this latter figure is calculated differs according to whether we are looking at wide or narrow proportionality.

However, I take this paper to be a starting point for discussion. There is much more to be said about what the best account of subjective proportionality will look like. I have flagged some of these issues, and some of my own arguments have been tentative. One important further issue is what the best account of necessity under uncertainty will look like, and how subjective proportionality and subjective necessity ought to interact. For now, I hope I have at the least shown that not all plausible views include the direct comparison between expected harm and expected good which it is natural to reach for, and that the issue of subjective proportionality is a complex and important one.