

Defensive Harm, Consent, and Intervention

(Note: This is a pre-publication (and slightly longer) draft of a paper that is forthcoming in *Philosophy and Public Affairs*. Please cite the published versions.)

Abstract: Many think that it would be wrong to defend an individual from attack if he competently and explicitly refuses defensive intervention. In this paper, I consider the extent to which the preferences of victims affect the permissibility of defending *groups* or *aggregates*. These cases are interesting and difficult because there is no straightforward sense in which a group can univocally consent to or refuse defensive intervention in the same way that an individual can. Among those who have considered this question, the dominant view is that that consent imposes only an extremely weak constraint on defending groups. I argue that this is mistaken and defend a much more robust requirement. Indeed, on the account that I develop there are cases in which the refusal of a *single* member of a victim is enough to make it impermissible to defend a much larger group. At the heart of my account is the idea that consent functions as an internal component of the broader requirement that defensive harms be proportionate: if a victim validly refuses defensive intervention, the fact that defence will benefit him cannot be used to justify harming innocents as the lesser-evil. An important implication of this view is that what constitutes sufficient consent from the members of a victim group will vary on a case-by-case basis.

1. Introduction

Contemporary discussions of the ethics of defensive harm typically focus on cases of self-defence, featuring two main characters: aggressors and their victims. The central question is to explain why victims are intuitively permitted to harm aggressors, given the usual prohibition on killing and injuring. The standard response is that, under certain conditions, individuals can render themselves *liable* to defensive harm and thereby lose their normal right not to be harmed. Debate centres on what the relevant conditions for liability consist in.

I will say almost nothing about liability here (in the cases I discuss I assume the aggressor's liability is uncontroversial). Instead, my topic is cases of *other*-defence. These feature an additional cast member: *third-party rescuers*, who are able to harm aggressors in defence of their victims. To date, the ethics of defending others has been relatively unexplored.¹ Of course, one explanation for this gap would be if the morality of other-

¹ The following are the only article-length discussions that I am aware of: Cécile Fabre, 'Permissible Rescue Killings', *Proceedings of the Aristotelian Society* 109 (2009), 149-164; Cécile Fabre, 'Mandatory Rescue Killings', *Journal of Political Philosophy* 15 (2007), 363-384; Seth Lazar, 'Authorisation and the Morality of War', *Australasian Journal of Philosophy* 94, No. 2 (2016), 211-226; Seth Lazar, 'Associative Duties and the Ethics of Killing in War', *Journal of Practical Ethics* 1, No.1 (2013), 6-51; Russell Christopher, 'Self-Defense and Defense of Others', *Philosophy and Public Affairs* 7, No.2, (1998), 123-141; Eduardo Rivera-Lopez, 'Puzzles on

defence were entirely determined by principles of self-defence. On this view, if it is permissible for a victim to inflict a certain harm on an aggressor, it follows that it is permissible for a third-party to do so in his defence.² If that's right, then there is no work for a theory of other-defence to do.³

One way to justify a theory of other-defence is by showing that there are *sui generis* constraints on defending others that do not apply to self-defence. In this paper, I explore one important candidate, endorsed by several theorists, which focuses on the *victim's will*.⁴ The constraint is usually characterised as follows:

Consent Requirement: Other-defence is morally justified only if those to be defended consent to defensive intervention.

Though this captures the central idea, the requirement is implausibly restrictive as stated. For one, it implies that it would be impermissible to defend someone who is unable to consent, such as someone being attacked while sleeping or unconscious.⁵ But we can easily revise the requirement to avoid such implications:

Consent Requirement*: Other-defence is morally justified only if those to be defended do not validly refuse defensive intervention.

Though rather unwieldy, this standard for consent should be familiar from medical contexts. For example, if a patient requires immediate treatment for a serious illness but is unconscious or otherwise unable to consent, doctors may proceed with treatment without

Defending Others from Aggression', *Law and Philosophy* 25, No.3 (2006), 377-386; George Fletcher, 'Defensive Force as an Act of Rescue', *Social Philosophy and Policy* 7, No. 2 (1990), 170-179.

² Provided that defender's defensive options are relevantly similar to the victim's i.e. the harm is the least harmful means available to them, using force doesn't involve disproportionately harming bystanders, etc. On this point, see Fabre 'Permissible Rescue Killings', pp.161-162.

³ This is one interpretation of Judith Jarvis Thomson's remark that "the permissibility of X killing Y goes hand-in-hand with the permissibility of Z killing Y in defence of X." Judith Jarvis Thomson, 'Self-Defense', *Philosophy and Public Affairs* 20, No.4 (1991), 283-310 at p.306.

⁴ Fabre, 'Permissible Rescue Killings', pp.159-160; Christopher Finlay, 'Legitimacy and Non-State Political Violence', *Journal of Political Philosophy* 18, No.3 (2010), 287-312 at pp. 290-295; Jeff McMahan, 'Humanitarian Intervention, Consent, and Proportionality' in N. Ann Davis, Richard Kershen, and Jeff McMahan (eds), *Ethics and Humanity: Themes from the Philosophy of Jonathan Glover* (Oxford: Oxford University Press, 2011), 44-72 at p.49; Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life* (Oxford: Oxford University Press, 2002), p.415; Victor Tadros, *The Ends of Harm: The Moral Foundations of Criminal Law* (Oxford: Oxford University Press, 2011), pp.295-296. Seth Lazar also endorses a (highly qualified) version of the requirement. Lazar, 'Authorisation and the Morality of War'.

⁵ On this point, see Fabre 'Permissible Rescue Killings', pp.159-160; Finlay, 'Legitimacy and Non-State Political Violence', p.292.

wronging the patient (provided she has not issued prior instructions not to be treated).⁶ For ease of exposition I will continue to refer to the ‘consent’ requirement on defensive harm, despite ‘the absence-of-valid-refusal requirement’ being the more accurate label.

In contrast to other standard components of a theory of defensive harm – such as liability, proportionality, and necessity – that focus on the normative situation between victims and aggressors, the consent requirement draws our attention to the relationship between victims and rescuers.⁷ It holds that it can be impermissible for a rescuer to harm an aggressor – and may even wrong *the beneficiary*⁸ – even if the rescuer would commit no wrong against the aggressor. If defensible, the consent requirement establishes a moral asymmetry between self- and other-defence: the fact that a victim is permitted to harm their aggressor does not entail that third-parties may do so, since the victim may refuse intervention.

The central question I want to focus on is how the consent-requirement governs large-scale cases of defensive force, in which *groups* or *aggregates* of individuals are the object of defence. In these cases, many find an analogue of the requirement very attractive. When it comes to military humanitarian intervention, for example, it is often argued that intervention is permissible only if carried out with the approval of the intended beneficiaries.⁹ But the point generalises to other contexts, such as the use of violence by non-state actors in defence of oppressed groups. The basic idea is that a belligerent group

⁶ This is roughly the position of the UK National Health Service. See <http://www.nhs.uk/conditions/consent-to-treatment/pages/introduction.aspx> [Last accessed 7th May 2017]

⁷ This is not to deny that consent can be relevant to the permissibility of defence in virtue of affecting the former relationship. For example, A may consent to being harmed by B and, by doing so, waive his right not to be harmed by B (for example, as part of a boxing match). Since B would not wrong A by harming him, it would be impermissible for a third-party to harm B in defence of A. In this paper I am exclusively concerned with cases in which victims refuse third-party intervention directly, without consenting to being harmed by their attackers.

⁸ Finlay ‘Legitimacy and Non-State Political Violence’, p.294; McMahan ‘Humanitarian Intervention, Consent, and Proportionality’, p.49.

⁹ See, for example, James Pattison, ‘Representativeness and Humanitarian Intervention’, *Journal of Social Philosophy* 38, No.4 (2007), 569-587; Fernando Tesón, ‘Eight Principles for Humanitarian Intervention’, *Journal of Military Ethics* 5, No.2 (2006), 93-113, at p.107; Fernando Tesón, ‘The Liberal Case for Humanitarian Intervention’ in J.L. Holzgrefe and Robert O. Keohane (Eds), *Humanitarian Intervention: Ethical, Legal, and Political Dilemmas* (Cambridge: Cambridge University Press, 2003), 93-130, at p.107; Jeff McMahan, ‘Just Cause for War’, *Ethics and International Affairs* 19, No.3 (2005), 1-21, at p.13; Eliav Lieblich, *International Law and Civil War: Intervention and Consent* (Abingdon: Routledge, 2013), p.256; Richard W. Miller, ‘Respectable Oppressors, Hypocritical Liberators: Morality, Intervention, and Reality’, in Deen K. Chatterjee and Don E. Scheid (eds), *Ethics and Foreign Intervention* (Cambridge: Cambridge University Press, 2008), 215-250, at p.224. Parallel claims are also sometimes made in discussions of non-military humanitarian interventions, such as aid operations. For example, Scott Wisor argues that “if a particular [aid] project is opposed by its intended beneficiaries, then this is a near-decisive reason to reject the distribution, even if the foreseeable consequences of the project are still very good.” Scott Wisor, ‘How Should INGOs Allocate Resources?’, *Ethics and Global Politics* 5, No.1 (2012), 27-48, at p.43.)

may only resort to rebellion, revolution, or insurgency if they have the support of the population they claim to be protecting.¹⁰ More broadly still, we might agree with Yitzhak Benbaji that in practically *any* defensive war, including wars of national defence, “groups on whose behalf the war is fought...are entitled to veto the war”.¹¹

However, despite its intuitive appeal, it is far from straightforward how we should apply the consent-requirement to these cases. The problem follows from an obvious disanalogy between individuals and groups. Collectives are not simply individuals ‘writ large’.¹² Rather, they are composed of individuals, each with a will and moral status of their own.¹³ Moreover, in any realistic case there is unlikely to be unanimity among the victims as to whether they desire defensive intervention. There is thus no straightforward sense in which ‘the victims’ or ‘the beneficiaries’ can consent to (or refuse) defensive intervention in the same way that an individual victim can. For example, imagine that State A is weighing up an intervention in State B, in order to prevent State B from committing serious and widespread right-violations against an ethnic group within its borders. What would it take for State A’s intervention to qualify as having the consent of the beneficiaries? What if 70% of the victims approve, but 30% protest? Or if 20% are in favour and 80% refuse? What if only a tiny handful of the beneficiaries want to be defended while the rest are opposed? Each of these distributions has a claim to being the correct interpretation of the consent-requirement, but how do we adjudicate between them? Until we settle this question, we cannot properly assess the case for intervention. Interestingly, the cases in which the consent requirement is most intuitive seem to be those in which it tells us the least.

Put in more general and precise terms, examples like this raise the question of how a theory of defensive harm should respond to what I call *multiple-victim* cases. These have the following features:

- (1) An aggressor threatens unjustified harm to several victims.

¹⁰ See Finlay, ‘Legitimacy and Non-State Political Violence’; Christopher Finlay, *Terrorism and the Right to Resist: A Theory of Just Revolutionary War* (Cambridge: Cambridge University Press, 2015), Ch.6; Lionel MacPherson, ‘Is Terrorism Distinctively Wrong?’ *Ethics* 117. No.3 (2007), 524-546; Michael Gross, *The Ethics of Insurgency* (Cambridge: Cambridge University Press, 2015), pp.37-44.

¹¹ Yitzhak Benbaji, ‘Legitimate Authority in War’, in Helen Frowe and Seth Lazar (eds), *Oxford Handbook of Ethics of War* (Oxford: Oxford University Press, forthcoming).

¹² As proponents of the ‘domestic analogy’ often imply. See, for example, Michael Walzer, *Just and Unjust Wars: A Moral Argument with Historical Illustrations*, 4th Ed (New York: Basic Books, 2006), Ch.4.

¹³ For a classic articulation of this point, see Charles Beitz, *Political Theory and International Relations* (Princeton: Princeton University Press, 1979), Ch.2. See also, Joel Feinberg, *The Moral Limits of the Criminal Law. Vol. 3: Harm to Self* (Oxford: Oxford University Press, 1986), pp.50-51.

- (2) A rescuer has the ability to prevent the threatened harm by imposing defensive harm on the aggressor.
- (3) Defending every victim is a condition of defending any.¹⁴
- (4) One sub-set of the members of the victim group consent to defensive intervention, another sub-set refuses.

In these cases, the consent requirement gives no guidance as to the permissibility of defence because it is indeterminate as to whether or not ‘those to be defended’ qualify as having consented (or, more accurately, as having not refused). To overcome this problem, the requirement needs to be supplemented with an *aggregation principle*, which tells us how we should move from the consent and refusal of individual victims to an all-things-considered judgement about whether or not the requirement has been met. My chief aim in this paper to provide such a principle.

This is no easy task. Allen Buchanan has recently expressed doubts that it is even possible, on the ground that “there apparently is no principled way of determining how widespread consent must be for the consent requirement to be satisfied”.¹⁵ Moreover, even if a non-arbitrary principle can be found, a different source of scepticism arises regarding whether the principle has any practical significance. This is the dominant view among those who have considered the aggregation question. According to these sceptics, the consent requirement imposes only a trivial constraint in multiple-victim cases, one that will be satisfied in practically any real-world example. More precisely, the claim is that the requirement is met just as long as *some* member(s) of the victim group do not validly refuse defensive intervention.

My central thesis is that this scepticism is mistaken, and that a much more robust aggregation principle governs multiple-victim cases. Indeed, according to the principle that I will defend, there are cases in which the refusal of a *single* victim can make it impermissible to defend a larger group, even if *every* other member explicitly consents. Moreover, I aim to show that we can get to this somewhat startling conclusion without abandoning the basic commitments that underpin the sceptical position.

¹⁴ Multiple-victim cases can be synchronic (in which several victims are threatened simultaneously) or diachronic (in which several victims are threatened consecutively). For ease of exposition, I discuss only synchronic cases, but the same story can be told for diachronic cases.

¹⁵Allen Buchanan, ‘Self-Determination, Revolution, and Intervention’, *Ethics* 126, No.2 (2016), 447-473 at n.10.

The paper proceeds as follows. In Section 2 I provide some intuitive support for the consent requirement in single-victim cases in order to motivate the question of whether, and to what extent, it constrains the permissibility of defending groups. In Section 3 I move on to cases featuring multiple victims and explain why many are attracted to the sceptical view. Its appeal lies in the plausible thought that requiring anything more robust than the consent of a single victim would involve giving group members an unacceptable degree of normative control over one another.

Section 4 presents the case for my robust aggregation principle. Drawing on discussions of the moral wrongness of paternalism, I argue that individuals have a normative power of control over whether others may act to promote their good. Just as others may not use my body and property without my authorisation, they may not justify their actions by appeal to reasons grounded in my good if I validly refuse to be benefitted. I term this the Power of Prudential Exclusion (PPE). In the context of defensive harm, the upshot of the PPE is that victims are able to constrain the reasons that rescuers can invoke in order to justify defensive intervention. If a victim validly refuses defensive intervention, the fact that defence will benefit him cannot be invoked by rescuers to justify harming innocents as the lesser-evil. When defending a group involves causing these harms – as in any realistic case of humanitarian intervention, revolution, or war – the refusal of victims serves to reduce the goods that are available to justify the use of defensive force. An important implication of this view is that what constitutes having ‘enough’ consent from the members of a victim group is context-sensitive. It will depend on (i) how large the group of beneficiaries is, (ii) how many victims validly refuse defensive intervention, and (iii) the extent to which innocents will be harmed in the course of defending the group. To return to our earlier question of whether State A’s intervention would satisfy the consent requirement if 70%, 20%, or 1% of the beneficiaries approve, my contention is that, until we flesh out these details, any of these distributions could be sufficient.

In Section 5 I consider an important objection to my proposal. The central thrust is that individual victims are subject to duties to rescue their co-victims, and that these duties place significant limits on their power to refuse defensive intervention. I argue that, contrary to first impressions, victims are not obligated to assist their co-victims by consenting to defensive harm in the relevant cases. In fact, they have no moral reason to do so *at all*. Their power to refuse therefore remains unconstrained. I conclude in Section 6 by demonstrating some of the wider practical implications of my account, focussing on the case of political protest.

2. Support for the Consent Requirement*

To help see the appeal of the consent requirement, consider the following single-victim case:

Elevator 1: Victim is taking a ride in an elevator. Aggressor begins to saw through the elevator cable because he hates Victim and wishes him dead. Rescuer is walking by and is able to shoot Aggressor with her sniper rifle, thereby saving Victim. However, Victim refuses intervention because he is deeply committed to non-violence.

In cases like this, many think that it would be morally impermissible for Rescuer to defend Victim (on the assumption that Victim has sufficient information, is not irrational, is not being coerced, etc.).¹⁶ The underlying idea is that victims occupy a privileged position within the morality of defensive harm. It's *the victim's* life that is at stake, not anybody else's, and so only he gets to decide if and how it is defended.¹⁷ This tracks the more general anti-paternalist thought that we may not interfere with a competent individual's self-regarding choices simply for the sake of her own good. On these matters, the individual is sovereign.¹⁸

¹⁶ Jeff McMahan, for example, finds it "intuitively obvious". 'Humanitarian Intervention, Consent, and Proportionality', p.49. See also the references in note 4 above.

¹⁷ Fabre, 'Permissible Rescue Killings' p.159.

¹⁸ It might be objected that appealing to anti-paternalist intuitions in order to support the consent-requirement is a mistake, because paradigmatic cases of paternalism have features that are lacking in cases like *Elevator 1*. For example, one might take paternalism to involve, (i) second-guessing the paternalisee's judgement about what is good for them, (ii) acting due to doubts about the paternalisee's having the required willpower to stick to their resolutions, and/or (iii) violating some independent claim of the paternalisee, such as their rights to bodily integrity or property. However, even if these are common features of paternalism, they are not plausibly necessary conditions, as the following two cases show:

Experiment: Benny freely and competently decides take part in a medical experiment that will make a major contribution to finding a cure for a serious genetic disease (which Benny does not have). The experiment has serious side-effects that will reduce Benny's life expectancy by one year. There is no compensation available for participants in the experiment. Out of concern for Benny's welfare, Sally locks Benny in his house so he misses the experiment.

Chocolate: Tony is an expert in nutrition who enjoys eating fatty foods, and chocolate in particular. He is fully aware that his diet is unhealthy and will cause him a premature death, but freely and competently decides that he prefers gustatory pleasures over an increased lifespan, and resolves to live a shorter life full of chocolatey bliss. Out of concern for Tony's welfare, Sally buys all the chocolate from every shop in Tony's hometown in order to stop him from purchasing any.

I take it that Sally clearly acts paternalistically (and impermissibly) in both these cases. But in *Experiment* there is no disagreement between Sally and Benny about whether his choice would be good for him. Participating in the experiment would obviously be bad for him, but he has chosen to sacrifice his welfare for non-prudential reasons (the good of others). So, Sally's benevolent interference can't involve second-guessing Benny's judgement about what is good for him (For related points, see Feinberg, *Harm to Self*, p.62; Seana Valentine Shiffrin, 'Paternalism, Unconscionability Doctrine, and Accommodation', *Philosophy and Public Affairs* 29, No.2 (2000), 205-250, at p.215; Tom Beauchamp, 'Paternalism and Bio-Behavioral Control', *The*

It is worth stressing, however, that we need not endorse a particularly strict form of anti-paternalism or self-sovereignty in order to accept the consent requirement in principle. Here I have in mind the Mill-inspired view that provided an individual's choice is sufficiently voluntary, benevolent interference is impermissible.¹⁹ While this would certainly support the consent requirement, it is not necessary. As formulated, the requirement holds that the absence of *morally valid* refusal is necessary for permissible defence. This is compatible with a range of views about the conditions under which refusal is valid, including those that place constraints on the *content* of choices, as well as the procedure by which they are made. For example, one might hold that individuals cannot validly refuse defensive intervention if they do so for reasons that are immoral, unimportant, inauthentic, or just plain silly.²⁰

To endorse the consent requirement, all one need accept is that (i) interference with at least *some* self-regarding choices for the sake of the chooser's good is prohibited, and (ii) that this includes at least *some* choices not to be defended by others. Of course, a complete account of the requirement will need to take a stand on what distinguishes valid from invalid refusal. But, for our purposes, we need not settle this here, since the question of determining the correct aggregation principle is independent of which particular package of procedural and content-based constraints one accepts. (In Section 5 I will consider a special kind of validity condition – neither procedural nor content-based – which is not independent in this way). Given these fairly minimal commitments, one need

Monist 60, No.1 (1976), 62-80, at p.78; Stephen Darwall, 'The Value of Autonomy and Autonomy of the Will', *Ethics* 116, No.2 (2006), 263-284; Scoccia, 'The Right to Autonomy and the Justification of Hard Paternalism', p.77.) In the *Chocolate* case, Sally is not motivated by any concerns about Tony's lack of willpower. Indeed, it is precisely because he will stick to his unhealthy resolution that she is moved to interfere. Furthermore, Sally violates no independent claim of Tony's by buying all the chocolate. If she were to purchase all the chocolate simply to consume it herself, or give it away to charity, Tony would have no legitimate complaint (For similar arguments against the 'independent claims' condition, see Shiffrin, 'Paternalism, Unconscionability Doctrine, and Accommodation', p.213; Feinberg, *Harm to Self*, p.9; Jonathan Quong, *Liberalism Without Perfection* (Oxford: Oxford University Press, 2011), pp.79-80.) Hence, the absence of features (i)-(iii) in cases like *Elevator 1* does not undermine the idea that Rescuer acts impermissibly in virtue of broadly anti-paternalist considerations. Thanks to an anonymous associate editor for prompting me to consider these points.

¹⁹ For a classic modern defence, see Feinberg, *Harm to Self*. See also, Richard Arneson, 'Mill versus Paternalism', *Ethics* 90, No.4 (1980), 470-489.

²⁰ See, for example, Richard Arneson, 'Joel Feinberg and the Justification of Hard Paternalism', *Legal Theory* 11 (2005), 259-284; Victor Tadros, 'Consent to Harm', *Current Legal Problems* 64 (2011), 23-49; Peter de Marneffe, 'Avoiding Paternalism', *Philosophy and Public Affairs* 34, No.1 (2006), 68-94; Peter de Marneffe, 'Self-Sovereignty and Paternalism' in Christian Coons and Michael Weber (eds), *Paternalism: Theory and Practice* (Cambridge: Cambridge University Press, 2013), 56-73; Danny Scoccia, 'Paternalism and Respect for Autonomy', *Ethics* 100, No.2 (1990), 318-334; Danny Scoccia, 'The Right to Autonomy and the Justification of Hard Paternalism' in Christian Coons and Michael Weber (eds), *Paternalism: Theory and Practice* (Cambridge: Cambridge University Press, 2013), 74-92; Daniel Brudney and John Lantos, 'Agency and Authenticity: Which Values Ground Patient Choice?', *Theoretical Medicine and Bioethics* 32, No.4 (2011), 217-227.

not share the intuition that defence would be impermissible in examples like *Elevator 1* in order to accept the consent requirement. This may just be a case in which one's preferred account of validity deems Victim's refusal to be morally ineffective.²¹

While most theorists who have considered cases like *Elevator 1* accept some version of the consent-requirement, Seth Lazar has recently offered a general argument against the idea that victims can veto third-party intervention.²² His main argument appeals to the fact that a victim's life and wellbeing is agent-neutrally valuable, giving all agents a moral reason to protect it. Since a victim "cannot by mere fiat prevent the continuation of his life from being an agent-neutral good", Lazar concludes that any suitably situated third-party may permissibly defend him.²³

²¹ Here are some slightly more complicated cases, in which the consent requirement seems most compelling (at least to me):

Excused Aggressor. Victim is taking a ride in an elevator. Aggressor begins to saw through the elevator cable, because he mistakenly (but not fully-culpably) believes that Victim will murder several innocent people unless Aggressor kills him. Rescuer is walking by and is able to shoot Aggressor with his sniper rifle, thereby saving Victim. However, Victim refuses intervention because he is deeply committed to non-violence.

Special Relationship. Victim is taking a ride in an elevator. Aggressor, who is Victim's child, begins to saw through the elevator cable because he hates Victim and wishes him dead. Rescuer is walking by and is able to shoot Aggressor with her sniper rifle, thereby saving Victim. However, Victim refuses intervention because he does not want his child to be killed for his sake.²¹

Non-Lethal Aggressor. Aggressor is about to kidnap Victim and lock him in his secret basement for six months because he hates Victim and wishes to deprive him of his liberty. Rescuer is walking by and is able to shoot Aggressor with her sniper rifle, thereby saving Victim. However, Victim refuses intervention because he is deeply committed to non-violence.

I assume that Aggressor remains liable to defensive killing in all three examples. But I take it that even the sceptical reader will feel the intuitive pull of the consent requirement in at least some of these cases. Moreover, the additional factors introduced above will be present in many real-world multiple-victim cases. For example, excused aggressors are fairly ubiquitous in armed conflicts, since many combatants who participate in unjust wars possess partial excuses for doing so, due to immaturity, ignorance, or duress. Similarly, non-lethal aggressors are present in cases of so-called 'bloodless invasion', in which the invader's aims only involve imposing non-lethal harms on their victims, such as territorial annexation and political repression (lethal force will only be employed if resistance is met.) Admittedly, the case of special relationships does not generalise so straightforwardly (though it would still clearly hold in cases where one family member attacks several others). But it may still plausibly apply in cases of *intra*-state violence, where oppressors and oppressed are often specially related to one another, albeit by non-familial bonds. So, even if victim refusal is only morally effective when these factors are in play, this does not significantly restrict the practical application of the consent-requirement, nor make the search for an aggregation principle any less pressing.

²² Lazar, 'Authorisation and the Morality of War', 216-218.

²³ *Ibid.*, p.217. Lazar accepts that something like the consent requirement may apply in a limited range of cases, in which killing is justified on the basis of agent-relative reasons, such as those grounded in special relationships. However, the existence of these justifications is highly contested.

However, this isn't a convincing argument, for three reasons. First, the view that it is impermissible to defend victims against their will need not rest on the claim that their refusal "negates the agent-neutral value of saving their life", as Lazar puts it.²⁴ Rather, the most natural interpretation of the idea is that refusal generates a *new* moral consideration against defending them that *defeats* the agent-neutral, welfare-based reasons in favour.²⁵

Second, there is simply no connection between an individual's continued life being agent-neutrally valuable and the question of whether that can veto defence. At base, the agent-neutral/agent-relative distinction concerns the *scope* of moral reasons; it is about whether or not a reason makes "some essential reference to the agent to whom they apply."²⁶ But formally classifying a reason as neutral or relative has no bearing on the substantive question of what other considerations – such as individuals' self-regarding choices – are able to defeat that reason. And that is the question at stake here. To illustrate: works of art may well have significant agent-neutral value, but this fact does not undermine the art owner's right to decide what happens to it.²⁷

Third, and most importantly, the argument from agent-neutrality has extremely counter-intuitive implications when applied to other cases. Consider the following example, familiar from the medical ethics literature:

Patient: Patient is ill and will die unless he receives a blood transfusion. Doctor can give the transfusion easily and safely. However, Patient explicitly and competently refuses medical intervention because having the treatment would go against his deep religious commitments.

Doctor would clearly act impermissibly if she gave Patient the transfusion against his will. This would be a paradigmatic case of wrongful paternalism.²⁸ But *Patient* shares some important features with *Elevator 1*. In each case, (i) a series of events will result in an individual's death unless prevented, (ii) the individual's continued life is agent-neutrally valuable, and (iii) the individual refuses third-party intervention that will save their life. If,

²⁴ Ibid., p.218.

²⁵ The notion of defeat here covers cases in which one reason *outweighs* another, as well as cases in which one reason defeats another by some other mechanism, such as by *exclusion* or *cancellation*. For this usage, see John Gardner and Timothy Macklem, 'Reasons', in Jules Coleman and Scott Shapiro (eds), *The Oxford Handbook of Jurisprudence and Philosophy of Law* (Oxford: Oxford University Press, 2002), 440-475, at pp.463-464.

²⁶ Ibid., citing David McNaughton and Piers Rawling's influential formulation of the distinction. David McNaughton and Piers Rawling, 'Value and Agent-Relative Reasons', *Utilitas* 7, No.1 (1995), 31-47.

²⁷ Thanks to an anonymous associate editor for this example.

²⁸ Revealingly, even defenders of 'hard' paternalism such as Danny Scoccia take it to be a serious objection to their view if it permitted interference in cases like this. Scoccia, 'The Right to Autonomy and the Justification of Hard Paternalism', at p.90.

as Lazar claims, the agent-neutral value of continued life ensures the permissibility of intervention in cases like *Elevator 1*, even in the face of competent refusal, then intervention should also be permissible in *Patient* (provided that it would better promote Patient's welfare than not intervening). In fact, intervention should be much easier to justify in *Patient* than in *Elevator 1* in one important respect: it doesn't involve intentional killing! Yet this is clearly the wrong result. More generally, the argument from agent-neutrality seems to entail the extreme conclusion that interference with individuals' choices is permitted whenever doing so overall promotes their wellbeing. Reflection on cases like *Patient* thus provides an additional source of support for the consent requirement.²⁹

I will postpone a deeper discussion of the consent requirement and its justification until Section 4. The points made so far are intended to demonstrate the initial plausibility of the requirement in single-victim cases, and thereby motivate the question of whether, and to what extent, the requirement constrains defence when multiple victims are threatened.

3. Scepticism About Consent in Multiple-Victim Cases

The dominant view among theorists who have considered these cases is that consent imposes little, if any, independent constraint on defensive harm. This scepticism is typically defended in opposition to a natural proposal (often made in discussions of humanitarian intervention and the use of force by non-state actors³⁰) which calls for a majoritarian aggregation principle. More specifically:

²⁹ One might try to resist this claim by pointing to differences between cases like *Patient* and *Elevator 1*. For example, one might argue that *Elevator 1* involves preventing *injustice* whereas *Patient* only involves preventing *misfortune*. In Section 4 I consider a variation on the case in which Patient is ill due to being poisoned by an aggressor. Since intervention remains intuitively impermissible in this case, it doesn't look like the injustice/misfortune distinction makes a difference to whether individuals may veto third-party intervention. Alternatively, one might argue that it makes a difference that intervention in *Patient*, but not in *Elevator 1*, involves violating an independent claim of the beneficiary's, such as rights to bodily integrity or property. However, this move implausibly makes big differences in permissibility turn on minor differences between cases. For example, if the force of the consent-requirement is limited to cases involving independent rights violations, this implies that intervention would be permissible in *Elevator 1*, yet impermissible in slight variations on the case in which defending Victim requires Rescuer to scratch Victim's arm (thereby violating his right to bodily integrity), or use Victim's sniper rifle (thereby violating his property rights).

³⁰ See, for example, Miller, 'Respectable Oppressors, Hypocritical Liberators: Morality, Intervention, and Reality', p.224. Eliav Lieblich, Christopher Finlay, Michael Gross, and Lionel MacPherson each understand the requirement in terms of a belligerent agent being 'representative' of the group they claim to defend. Some form of majoritarianism is a natural (though not unique) interpretation of this idea, and many of the examples these authors use suggest they have something like this in mind. Lieblich, *International Law and Civil War*, p.256; Finlay, 'Legitimacy and Non-State Political Violence'; MacPherson, 'Is Terrorism Distinctively Wrong?'; Gross, *The Ethics of Insurgency*, pp.37-44.

The Majority Consent Requirement: Other-defence is morally justified only if a majority (or super-majority) of the members of the victim group consent to defensive intervention (or do not validly refuse).³¹

This view has obvious appeal. After all, few dispute that democracy is a good thing. In many contexts majority rule seems the appropriate way of deciding how to proceed given disagreement. The problem, however, is that when it comes to defending basic interests – such as life and limb – the majoritarian proposal is far from intuitive. To demonstrate, consider the following case, based on an example of Andrew Altman and Christopher Heath Wellman’s³²:

Elevator 2: Five Victims are taking a ride in an elevator. Aggressor begins to saw through the elevator cable because he hates Victims and wishes them dead. Rescuer is walking by and is able to shoot Aggressor with her sniper rifle, thereby saving the Victims. However, while one Victim consents to Defender’s intervention, the four remaining Victims refuse because of their deep commitment to non-violence.

Here defence seems clearly permissible, despite the fact that 80% of the victims explicitly refuse. Moreover, the intuition persists even if we specify that the victims share a special relationship, such as family membership. These verdicts track the powerful idea that individuals’ basic interests place limits on majority rule.³³ When these interests are threatened, individuals may have the right to decide *for themselves* whether they are defended, but not to decide for others.³⁴ As Altman and Wellman put it, “It seems dubious to hold that a group has this type of normative dominion over its members.”³⁵ As in the case of *Elevator 1*, the value of individual sovereignty seems to ground the intuitive verdict.

³¹ For discussion of several distinct majoritarian proposals, albeit in the context of non-defensive interventions, see Kalle Grill, ‘Liberalism, Altruism, and Group Consent’, *Public Health Ethics* 2, No.2 (2009), 146-157.

³² Andrew Altman and Christopher Heath Wellman, ‘From Humanitarian Intervention to Assassination: Human Rights and Political Violence’, *Ethics* 118, No.2 (2009), 228-257, at p.244.

³³ For a classic articulation, see Brian Barry, ‘Is Democracy Special’, in Peter Laslett and James Fishkin (ed) *Philosophy, Politics and Society: Fifth Series* (Yale University Press, 1979), 155-196, at pp.170-171.

³⁴ This leaves open the possibility that when other kinds of interests are threatened a majoritarian consent requirement could be appropriate. For discussion, see [Reference Removed].

³⁵ Altman and Wellman, ‘From Humanitarian Intervention to Assassination’, p.243. For similar arguments, see McMahan, ‘Humanitarian Intervention, Consent, and Proportionality’, p.52; Helen Frowe, ‘Judging Armed Humanitarian Intervention’, in Don E. Scheid (ed), *The Ethics of Armed Humanitarian Intervention* (Cambridge: Cambridge University Press, 2014), 93-110 at pp.108-109; Tesón, ‘The Liberal Case for Humanitarian Intervention’, p.106.

We can draw two sceptical conclusions from cases like *Elevator 2*, one stronger than the other. The weaker conclusion is that majoritarian aggregation principles are false: it is not true that the consent (or non-refusal) of the majority is a necessary condition for permissibly defending groups of victims from serious threats. However, sceptics about the consent-requirement go further than this, and endorse the stronger conclusion that refusal can only render defence impermissible when it is *unanimous* among the members of the victim group. To support this view, consider variations on *Elevator 2*, in which increasing numbers of refusing victims are present in the elevator. It still seems clear (at least to me) that defence would be permissible even if the refusers greatly outnumber the single consentor.³⁶ Putting things more precisely, we can interpret the sceptics as proposing the following, extremely weak, aggregation principle:

The Minimal Consent Requirement: Other-defence is morally justified only if *some* member(s) of the victim group consent(s) to defensive intervention (or do not validly refuse.)

On this view, while the consent requirement applies *in principle* in multiple-victim cases, there will be practically no instances in which a lack of consent *in fact* renders defence unjustified.³⁷ This has important practical implications. For example, if a military humanitarian intervention meets all the standard conditions of just war (just cause, proportionality, last resort, etc.) then even very widespread refusal among the victim population cannot itself render the intervention impermissible. As Jeff McMahan, one of the chief consent-sceptics, puts it:

“that the familiar principles of *jus ad bellum*, and especially the proportionality condition, do all the substantive work that the requirement of consent is supposed to do is largely correct as an account of the conditions of *objectively* justified humanitarian intervention.”³⁸

³⁶ Fabre suggests that she also shares the sceptical view, noting that in a diachronic multiple-victim case, an individual victim’s refusal “would not carry much weight if any”. ‘Fabre, Permissible Rescue Killings’, n.11.

³⁷ As McMahan and Frowe each point out, in moderately large victim groups there will be invariably be members who are not capable of *valid* refusal, such as young children. Hence, the Minimal Consent Requirement is likely to be met even in cases where every victim explicitly refuses. McMahan, ‘Humanitarian Intervention, Consent, and Proportionality’, p.53. Frowe, ‘Judging Armed Humanitarian Intervention’, p.109.

³⁸ McMahan, ‘Humanitarian Intervention, Consent, and Proportionality’, p.52. Frowe similarly concludes that “we should not include consent as a separate component of justification for intervention.” Frowe, ‘Judging Armed Humanitarian Intervention’, p.109.

And, more strongly:

“[consent] has no role among the conditions of objective justification in ideal theory”³⁹

This is consistent with consent playing an important *derivative* role in justifying other-defence. For example, widespread refusal may provide good evidence that defence would not be justified on some other ground, or render intervention unlikely to succeed (for example, because success requires the cooperation of the victims).⁴⁰ The sceptics’ position is that consent imposes only a trivial *independent* constraint on intervention in multiple-victim cases.

4. In Defence of Consent

In what follows I provide a defence of the consent requirement against this scepticism. As my starting points, I accept the intuitive verdicts in both *Elevator 1* and *Elevator 2* and the sceptics’ rejection of the Majority Consent Requirement. I also accept the deeper liberal/individualist picture of morality that underpins – either explicitly or implicitly – the sceptical position. But I will deny the Minimal Consent Requirement follows from these judgements. Instead, we should endorse a much stronger aggregation principle, according to which defending a group can be impermissible in virtue of a *single* victim refusing defensive intervention.

4.1 The power to control one’s good

I start by providing a deeper moral basis for the consent requirement, that provides a unified explanation of the contrasting verdicts in *Elevator 1* and *Elevator 2*.

Here is the basic idea. When it comes to our bodies and property, each of us has the moral power to exclude others from using them, simply because they are *ours* and we

³⁹ Ibid., p.53. McMahan does concede that consent may matter under conditions of epistemic uncertainty. I will not consider this idea here, since my aim is to provide a defence of the consent requirement at the objective or ‘fact-relative’ level.

⁴⁰ Altman and Wellman, ‘From Humanitarian Intervention to Assassination’, pp.244-245; Allen Buchanan, ‘The Ethics of Revolution and Its Implications for the Ethics of Intervention’, *Philosophy and Public Affairs* 41, No.4 (2013), 292-323; Pattison, ‘Representativeness and Humanitarian Intervention’, pp.580-583; Benbaji, ‘Legitimate Authority in War’.

are in charge of them. Others may only use our material resources if we authorise them to do so and we are wronged if they do so without our consent, even if they improve those resources by using them.⁴¹ My contention is that, in addition to these more prosaic rights, each of us has structurally similar powers of control regarding whether, and for what purposes, others may use our *good*. One person uses another's good when she justifies her actions by appeal to the fact that it will promote the other's welfare. On this view, reasons pertaining to an individual's good can be thought of as a kind of resource – a *moral* resource. And just as I can exclude you from using my body or property by withholding my consent, I have the power to prevent you from justifying your actions by appealing to the fact that you will benefit me, by repudiating those benefits. By exercising this power, I render my good morally inadmissible for the purposes of justifying your action, and you commit a form of trespass against me if you attempt to do so. For want of a convenient label, term this normative ability the *Power of Prudential Exclusion* (PPE).⁴² An important upshot of the PPE is that individuals can, by controlling access to justifying reasons, render actions impermissible that would otherwise be justified. As I will explain below, it is this consequence of our self-regarding powers that gives the consent requirement its moral foundations.

But some further work needs to be done first. What I have said so far simply sets out the structural features of the PPE. But why should we think that individuals have this power? And what grounds or justifies it?

In my view, the most compelling evidence for the PPE comes from noting the central role it plays – either explicitly or implicitly – in prominent accounts of paternalism's distinctive wrongness, as well as in explaining judgements about particular cases.⁴³ I will not attempt to offer a definition of paternalism.⁴⁴ Instead, I will (rather roughly) take

⁴¹ For a helpful discussion of 'use-based-wrongs', see Arthur Ripstein, *Private Wrongs*, (Cambridge, Mass: Harvard University Press, 2016), Ch.2.

⁴² While preparing the final version of this paper, I became aware of a very similar proposal defended in unpublished work by Seth Lazar. Lazar also argues that your rights of self-ownership extend to your interests, and that one important upshot of these rights is that "whether those interests count for or against some action is, at least in part, to be determined by your choice." However, Lazar deploys this idea for different purpose to mine. Whereas I focus on the implications of our rights over our good for how *others* may permissibly act, Lazar is concerned with how these rights affect the deontic status of *our own* actions. In particular, Lazar argues that these ownership rights ground agent-centred options to favour and sacrifice one's own interests, even when doing so does not promote aggregate wellbeing. Seth Lazar, 'Self-Ownership and Agent-Centred Options' (unpublished manuscript). I briefly discuss some of Lazar's arguments for the general right below.

⁴³ If Lazar is correct that positing the right to control how one's interests may be used grounds central features of commonsense morality, then this provides another source of support for something like the PPE.

⁴⁴ For a detailed overview of debates regarding the definition of paternalism, see Gerald Dworkin, 'Defining Paternalism' in Christian Coons and Michael Weber (eds), *Paternalism: Theory and Practice* (Cambridge:

central cases of paternalism to involve: (i) Some form of interference with an individual, (ii) that is carried out against her will, and (iii) for the sake of her own good. On a fairly standard conception, paternalism is wrong (when it's wrong) because the paternalisee's autonomy has priority over considerations of her wellbeing or good. The question is how to characterise this priority. According to one influential view, we should understand it in terms of competent individuals having *authority* with respect to their good.⁴⁵ When it comes to matters concerning her own welfare, the individual is 'the boss' and her will settles the question of whether others may act to promote her good. The paternalistic wrong lies in denying individuals rightful authority over themselves. By contrast, paternalism may be permissible when one's self-regarding authority is compromised in some way (for example, due to duress, deception, or lack of capacity).

I think proponents of authority-accounts of paternalism are appealing to something very similar to the PPE. This is made explicit in Daniel Groll's recent articulation of the view. For Groll, my having authority with respect to my good consists in having the normative power to give others exclusionary reasons: that is, second order-reasons not to act for certain first-order reasons.⁴⁶ My self-regarding choice "is meant to supplant the reason-giving force of other considerations not because it outweighs those other considerations but because it is meant to silence or exclude those other considerations."⁴⁷ Chief among the excluded reasons are those grounded in the fact that treating me contrary to my will is likely to promote my welfare. We need not endorse Groll's specific exclusionary-reasons treatment of authority in order to appreciate its affinities with the PPE. The key point is that, as a result of exercising my will, certain reasons – those pertaining to my good or welfare – no longer contribute to justifying others' actions. We can remain neutral on whether the relevant normative mechanism is

Cambridge University Press, 2013), 25-38. See also, Emma C. Bullock. 'A Normatively Neutral Definition of Paternalism', *Philosophical Quarterly* 65, No.258 (2015), 1-21.

⁴⁵ This view has been developed most thoroughly by Seana Shiffrin and Daniel Groll, but it seems implicit in many discussions of paternalism. On Shiffrin view, many paradigmatic cases of paternalism involve "a person's aiming to take over or control what is properly within the agent's own legitimate domain of judgement or action." ('Paternalism, Unconscionability, and Accommodation', p216). On Groll's view, a competent individual is the "de jure ultimate decision-maker" with respect to matters that concern her good ('Paternalism, Respect, and the Will', *Ethics* 122, No.4 (2012), 692-720, at p.700). For another endorsement, see Cholbi, 'Paternalism and Our Rational Powers' *Mind* 126, No.501 (2017), 123-153. Joel Feinberg's treatment of how paternalism violates autonomy also suggests he is an advocate of the authority-view (see, for example, Feinberg, *Harm to Self*, p.52).

⁴⁶ For seminal discussion, see Joseph Raz, *Practical Reason and Norms* 2nd Ed. (Oxford: Oxford University Press, 1999). David Enoch also places the notion of an exclusionary reason at the heart of his account of the wrong of paternalism. David Enoch, 'What's Wrong with Paternalism: Autonomy, Belief, and Action', *Proceedings of the Aristotelian Society* 116, No.1 (2016), 21-48.

⁴⁷ Groll, 'Paternalism, Respect, and the Will', p.701.

best cashed out in terms of exclusion, cancelation, silencing, nullification, invalidation, or some other notion.

An important virtue of authority-accounts of paternalism – and therefore the PPE – is that they enable us to explain the intuitive difference between cases like the following⁴⁸:

Medicine 1: Bob is ill and will die unless he takes some medicine that is delivered via a gas. However, he competently refuses medical treatment because the medicine contains a substance that his religious convictions prohibit ingesting. Doctor pumps the gas into Bob’s hospital room in order to save his life. Bob breathes the gas.

Medicine 2: Dave is ill and will imminently develop a lifelong debilitating condition unless he takes some medicine that is delivered via a gas. He consents to Doctor giving him the medicine. However, due to time constraints, the only way to treat him is to immediately pump the gas into Dave’s hospital ward. Bob is the hospital bed next to Dave recovering from a slipped disc, which has rendered him immobile. Bob refuses to be given the medicine because the medicine contains a substance that his religious convictions prohibit ingesting. The cost to Bob of breathing the gas is sufficiently high that he is not under a duty to bear it in order to save Dave the cost of the debilitating disease. Doctor pumps the gas into the ward in order to save Dave from the debilitating condition. Both Bob and Dave breath the gas.

I assume it would clearly be wrongfully paternalistic for Doctor to force Bob to consume the gas against his will in *Medicine 1* (this is simply the *Patient* case with slightly different apparatus). In *Medicine 2*, by contrast, it would not be paternalistic, and may well be morally permissible for Doctor to do so (at very least, there seems to be a significant moral difference between the two cases). This comparison reveals, interestingly, that it can be permissible to transgress an individual’s rights to autonomy and bodily integrity in order to promote the good (preventing the harm of a debilitating disease), yet impermissible to violate those very same rights in order bring about a *greater* good (preventing the harm of death). The PPE neatly explains how this can be so. What matters is not simply *how much* wellbeing is promoted, but *whose* wellbeing it is and whether they have exercised their powers of control over it. While we may not transgress Bob’s rights for the sake of *Bob’s* good (because he has excluded us from appealing to those reasons), we may do so for the sake of Dave’s *lesser* good because those reasons remain ‘in play’ and may be sufficient to justifying overriding Bob’s rights as the lesser-evil.

⁴⁸ For a range of other cases in which authority-accounts are well-placed to capture important distinctions, see Groll, ‘Paternalism, Respect, and the Will’.

Moving on to the justification of the PPE. I believe the PPE can be grounded in a plurality of considerations. A complete account of these factors is beyond the scope of this paper. Instead, let me briefly sketch five promising candidates.

First, and most obviously, the PPE responds to the fact that we are intimately and uniquely connected to our own good, perhaps even more so than our own bodies. As Seth Lazar argues, we are not only affected by the satisfaction and frustration of our interests as no one else can be, they also play a central role in constituting our identity: there is a sense in which we just *are* our interests.⁴⁹ Given the degree to which we are enmeshed with our own wellbeing, it seems fitting that we have rights to control how others may use it, just as we have rights over what others may do with our bodies.

Second, the PPE serves our interest in autonomy, in at least two ways. First, our having the ability to prevent others from acting to promote our good enhances our ability to shape our lives in accordance with values and projects that we endorse (autonomy as *self-determination* or *authenticity*). Second, by giving us control over the reasons for which others may act, the PPE protects our interest in not being controlled by others (autonomy as *independence*). In each case, the PPE increases the degree to which *what happens to us* depends on our choices.⁵⁰

Third, we plausibly have interests in being able to restrict access to our good independently of whether it helps us control others' behaviour. Common-sense morality contains a stringent constraint on using others for the sake of goals they don't share.⁵¹ What counts as using persons is typically taken to cover using their bodies and property. However, given the intimate connection between one's self and one's good, it seems reasonable to include the use of another's good within the scope of the constraint. The PPE is thus valuable because it gives us a measure of protection against being wrongfully used *as an end* in this way. In cases where I don't endorse others' plans to benefit me, this power enables me to morally divest myself from those plans, and prevent my good from justifying goals that I do not share. The PPE thus serves an interest in having *normative* control over the reasons that others may act on, separate from our interests in *de facto* control over what others do.⁵²

⁴⁹ Lazar, 'Self-Ownership and Agent-Centred Options'.

⁵⁰ On these two dimensions of autonomy, and how they can come apart, see Seana Valentine Shiffrin, 'Autonomy, Beneficence, and the Permanently Demented', in Justine Burley (ed), *Dworkin and His Critics* (Oxford: Blackwell Publishing, 2004), 195-217. See also Tadros, *Wrongs and Crimes*, Chs.8&11.

⁵¹ See, for example, Warren Quinn, 'Actions, Intentions and Consequences: The Doctrine of Doing and Allowing', *The Philosophical Review* 98. No.3 (1989), 287-312.

⁵² On normative interests, see David Owens, *Shaping the Normative Landscape* (Oxford: Oxford University Press, 2012).

Fourth, the PPE can be justified by appeal to the valuable relationships that it makes possible. By having the power to control access to my good, I am able to pick out certain persons as those who can treat me benevolently without violating a duty. This permission enables a form of intimacy and exclusivity that is integral to creating and sustaining many kinds of valuable relationship, such as those between friends and family members, and between care-givers and those cared for. Were we to lack the PPE, and our good were accessible to all, this valuable intimacy would not be possible.⁵³

Fifth, the PPE may not only serve our interests, but also partly constitute our moral *status*. According to a familiar strain of non-consequentialist thought, persons have value as ends in themselves, and not just in virtue of their amount of wellbeing their lives contain. Were we to matter only in the latter sense, we would be mere ‘receptacles’ of value and morality would permit imposing costs on us for the sake of net gains in aggregate wellbeing. In order to properly reflect the fundamental independence of persons, morality must contain constraints on promoting the good (such as rights against harm).⁵⁴

However, incorporating constraints only partly responds to the separateness-of-persons/value-receptacle worry. In addition, we need something like the PPE. To help see this, recall the *Medicine 1* and *Medicine 2* cases discussed above. In each case, Bob is equally protected from being sacrificed for the sake of the greater good: he has (non-absolute) rights that others not infringe his autonomy and bodily integrity. In this respect, morality accommodates Bob’s unique perspective with respect to his own life and projects. But under a moral system that lacked the PPE, there would be no deep difference between infringing Bob’s rights for the sake of the greater good *when the good is constituted by Bob’s wellbeing* (as in *Medicine 1*) and *when it is constituted by someone else’s* (as in *Medicine 2*). Bob would be treated as an interchangeable container of value, in precisely the manner ruled out by a concern for the separateness-of-persons.⁵⁵ Our moral status thus grounds the PPE.

4.2 Explaining the verdict in *Elevator 1*

⁵³ This is just a snapshot of a more general form of argument developed in great depth in Owens, *Shaping the Normative Landscape*. See also Tadros, *Wrongs and Crimes*, Ch.11 for helpful discussion.

⁵⁴ On this point, see Thomas Nagel, ‘Personal Rights and Public Space’, *Philosophy and Public Affairs* 24, No.2 (1995), 83-107; Frances Kamm, ‘Non-Consequentialism, the Person as an End in Itself, and the Significance of Status’, *Philosophy and Public Affairs* 21, No.4 (1992), 354-389.

⁵⁵ I am thus in agreement with Michael Cholbi’s suggestion that refraining from paternalism “seems to consist in honoring boundaries between persons”. Michael Cholbi, ‘Paternalism and Our Rational Powers’, *Mind* 126, No.501 (2017), 123-153, at p.145.

With the positive case for the PPE now in place, we can now return to the consent requirement. The idea that individuals have a power of exclusion over their good provides a deeper explanation of the intuitive verdict in *Elevator 1*.

Though it is often overlooked, the justification for defensively killing aggressors has two basic parts. Firstly, and most obviously, we appeal to the aggressor's liability. Liability functions negatively in the justification, by removing the normal constraint on intentional killing. But this is not sufficient to justify killing the aggressor. To say that someone is liable to defensive killing is just to say they would not be wronged by being killed for a particular defensive purpose (that they lack a claim-right against such treatment.) But the fact that an action would not wrong someone does not generate any positive reason to perform it, and we still have weighty moral reasons not to kill even liable individuals. Put simply: liable persons are still persons, whose lives and interests matter morally. Our duties of beneficence require us to give these facts considerable weight in our deliberations.⁵⁶ To see this, consider the intuitive constraints on defensive harm encapsulated in the requirement of necessity. If a victim can avoid being killed by a culpable aggressor by retreating at very little cost, or can avoid being killed by a gang of culpable aggressors by killing one aggressor rather than all, then they are morally required to take the least harmful option. Since, by hypothesis, the aggressors are all fully liable to defensive killing, this common-sense constraint shows that there are significant moral reasons against killing liable persons.

The justification for killing aggressors thus requires a second component, which identifies a positive moral reason in favour of killing. This reason is straightforward: killing the aggressor is necessary to protect the continued life and wellbeing of the victim. The obviousness of this explains why we tend to treat the justification for killing aggressors as exhausted by considerations of liability.

Once we recognise this additional welfarist component, we can see more precisely why defensively killing aggressors in cases like *Elevator 1* is impermissible. Victim's refusal prevents the welfarist component of the justification for defensive killing from being satisfied. By exercising his power of exclusion, Victim prevents a class of reasons from

⁵⁶ For elaboration on these sorts of reasons, see Joanna Mary Firth and Jonathan Quong, 'Necessity, Moral Liability, and Defensive Harm', *Law and Philosophy* 31, No.6 (2012), 673-701. See also, David Rodin, 'The Lesser Evil Obligation', in Saba Bazargan-Forward and Samuel C. Rickless (eds), *The Ethics of War: Essays* (Oxford: Oxford University Press, 2017), 28-45; Michele Madden Dempsey, 'Against Liability: A Reason-Based Account of Self-Defense', in Christian Coons and Michael Weber (eds), *The Ethics of Self-Defense* (Oxford: Oxford University Press, 2016), 211-232; Adil Ahmed Haque, *Law and Morality at War* (Oxford: Oxford University Press, 2016), p.120.

contributing to the justification of killing Aggressor: those grounded in the fact that he would be benefitted. As a result, there is insufficient *admissible* moral reason that Rescuer can invoke to justify defending Victim. Rescuer thus both lacks a justification for killing Aggressor and commits an additional paternalist wrong against Victim if they do so, by using a moral resource that Victim has authoritatively declared may not be used for that purpose.

One might object that this justification of the consent requirement is undermined by a familiar argument from more general debates about the permissibility of paternalism.⁵⁷ When confronted with a putatively paternalistic policy that seems both sensible and desirable, a common move in the anti-paternalist's repertoire is to argue that the policy can in fact be supported by non-paternalistic reasons.⁵⁸ For example, Seana Shiffrin famously argues that the state need not act paternalistically if it refuses to enforce contracts it deems to be exploitative or deeply unfair, because it need not be guided by the aim of protecting individuals from the harmful consequences of their voluntary choices. Instead, the state may permissibly refuse because it has a legitimate interest in avoiding complicity in exploitation.⁵⁹

A structurally similar line of argument may be advanced against my defence of the consent requirement. On this view, Rescuer need not act paternalistically by defending Victim in *Elevator 1* because there are alternative reasons that support killing Aggressor besides promoting Victim's good, and Rescuer may be acting solely for those reasons. Most obviously, Rescuer may be guided by the aim of preventing a serious injustice. Hence, we can grant victims the right to exclude rescuers from acting for their good, but deny that exercising this power renders other-defence impermissible.

In response, the first thing to note is that this objection relies on the assumption that victims of aggression lack control over the justice-based reasons for harming aggressors. But this assumption can be contested. As Cécile Fabre has argued in response to a similar objection, "it is not enough to ensure that justice obtains; one must also ensure that the right person [the victim] decides whether or not to bring about justice."⁶⁰ If Fabre is right that justice is sensitive to victims' preferences in this way, then the appeal to justice-based reasons for harming need not undermine the consent requirement. At very least, the

⁵⁷ Thanks to an anonymous editor for pressing me to consider this objection.

⁵⁸ Peter de Marneffe labels this 'the reconciliation strategy'. Peter de Marneffe, 'Avoiding Paternalism'.

⁵⁹ Shiffrin, 'Paternalism, Unconscionability Doctrine, and Accommodation'.

⁶⁰ Fabre, 'Permissible Rescue Killings', p.159. My parenthesis.

objector owes us an argument for why justice is preference-independent in the relevant sense. In what follows I will assume that such an argument can be given.

I grant that the objection may well succeed in showing that defending victims against their will does not always involve acting paternalistically. It seems plausible that paternalistic wrongdoing requires being motivated by the paternalisee's good, and Rescuers need not have that motivation. However, I highly doubt that it shows that defending victims against their will is *morally justified*. And this is what needs to be demonstrated in order to challenge my account of the consent requirement. The crucial point is that, on my view, the impermissibility of defending Victim is not explained by his being paternalistically wronged. Rather, it is explained by the fact that Victim – by exercising his PPE – prevents Rescuer from having a justifying reason for killing Aggressor (that Victim's good will be promoted). So, even if Rescuer is not motivated by Victim's good – and therefore does not paternalistically wrong him – she does not thereby acquire a justification for killing Aggressor, since Victim's good remains inadmissible. So, in order to show that Rescuer is justified in defending Victim, it is not enough to identify some alternative reasons for killing Aggressor, such as preventing injustice. It must also be shown, much more strongly, that those reasons are *sufficient* to justify intentionally killing Aggressor, independently of considerations of Victim's welfare.

While we may agree with Shiffrin that our reasons to avoid complicity in exploitation may be sufficiently important to justify failing to enforce an unfair contract, it seems implausible that our reasons to prevent injustice could be weighty enough to justify killing, even if the target is liable. Persons' lives and welfare matter a great deal, independently of whether those interests are protected by rights. This seems especially clear in cases where the aggressor is not fully culpable. My suspicion is that whatever intuitive force this idea has comes from the fact that preventing serious injustice and promoting welfare are typically practically inseparable. But the objection under consideration requires that preventing abstract, impersonal injustice *itself* is sufficient to justify killing, even if it were not good *for* anyone. That is a pretty extreme view.

To help see the relative weakness of justice-based reasons in life-or-death contexts, consider the following example:

Rock Rescue: Betty is taking a walk at the bottom of some cliffs, when she spots two rocks falling from the cliff tops. The first rock has been blown by the wind and will land on two innocent persons, killing one and breaking the other's leg. The second rock has been deliberately pushed by Villain and will kill

a third innocent person. Luckily, Betty has her ray gun with her and has time to vaporise one rock before it lands, but not both.⁶¹

If our reasons to prevent injustice were weighty enough to justify killing, then we should expect to judge that Betty ought to vaporise the second rock and prevent the murder. But it seems clear that instead she ought to vaporise the first rock and prevent the moderately greater harm.⁶² As Jeff McMahan puts it, “it would be perverse to prevent the murder rather than the accidental death together with the nonlethal injury.”⁶³ Moreover, even if each boulder only threatened a single life, it is still not obvious that preventing injustice would be important enough justify vaporising the second rock, rather than tossing a coin and giving each victim an equal chance of survival.⁶⁴ If our reasons to prevent injustice yield so easily to considerations of welfare and fairness, this strongly suggests that they are insufficiently weighty to justify killing. The justification for preventing unjust aggression is largely grounded in the importance of preventing harm, rather than preventing wrongs.⁶⁵

We can cast further doubt on the significance of preventing injustice by considering a variation on the *Patient* case discussed earlier:

Poisoned Patient: Patient has been poisoned by a malicious enemy, and will die unless he receives a blood transfusion. However, Patient explicitly and competently refuses medical intervention because having the treatment would go against his deep religious commitments. Doctor can give the transfusion easily, safely, and without Patient ever knowing.

In the original *Patient* case, it would be straightforwardly impermissible for Doctor to treat Patient against his will, as well as clearly paternalistic (since Doctor would justify his actions

⁶¹ Based on a case discussed, independently, by Jeff McMahan and Peter Singer, who draw very similar conclusions. McMahan, ‘Humanitarian Intervention, Consent, and Proportionality’, p.60; Peter Singer, ‘Bystanders to Poverty’, in N. Ann Davis, Richard Kershen, and Jeff McMahan (eds), *Ethics and Humanity: Themes from the Philosophy of Jonathan Glover* (Oxford: Oxford University Press, 2011), 187-201 at pp.195-198.

⁶² Frances Kamm agrees with the verdict in this type of case, but reports the intuition that preventing lesser unjust harms may be permissible in cases where one would prevent the unjust *act* and not just its harmful *effects*. However, this has the rather counterintuitive implication that the permissibility of Betty’s choice in *Rock Rescue* depends on whether she also has the option of preventing the second boulder from hitting its target by vaporising the stick that Villain would otherwise use to perform the act of pushing the rock off the cliff. Frances Kamm, ‘Self-Defence, Resistance, and Suicide’, in Helen Frowe and Gerald Lang (eds), *How We Fight* (Oxford: Oxford University Press, 2014), 75-86, at pp.75-76.

⁶³ McMahan, ‘Humanitarian Intervention, Consent, and Proportionality’, p.60.

⁶⁴ Derek Parfit, Victor Tadros, Liam Murphy, and Kieran Oberman also endorse the view that preventing injustice makes little or no difference in cases of equal harm. Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1986), p.48; Tadros *The Ends of Harm*, pp.105-106; Tadros *Wrongs and Crimes*, pp. 162-166; Liam B. Murphy, *Moral Demands in Non-Ideal Theory* (Oxford: Oxford University Press, 2000), pp.126-127; Kieran Oberman, ‘War and Poverty’ (unpublished manuscript).

⁶⁵ See Tadros, *The Ends of Harm*, p.106.

by appeal to Patient's good). In *Poisoned Patient*, by contrast, Doctor can appeal to the aim of preventing a serious injustice. But this doesn't seem to make the intervention much easier to justify (even if it absolves Doctor of the charge of paternalism).⁶⁶ But if reasons to prevent injustice are insufficient to affect permissibility of *saving* a life against the beneficiary's will, it is hard to believe that they are important enough to justify *killing* in contravention of the beneficiary's wishes, even if the target is liable.

There are, however, a limited class of cases in which non-victim-centric reasons may be sufficient to justify harming aggressors, and so permit third-parties to defend victims even in the face of valid refusal. These are cases in which a serious threat to a victim can be averted by imposing a relatively minor harm on the aggressor.⁶⁷ For example, if a rescuer can prevent an aggressor from killing her victim by spraining the aggressor's wrist, then it doesn't seem implausible that this harm could be justified by appeal to the value of preventing serious injustice, or even the benefit *to the aggressor* of not committing a serious moral wrong.⁶⁸ If so, then spraining the aggressor's wrist may be permissible even if the victim validly refuses defensive intervention, since the justification would not require appealing to the victim's good. Since it would be a troubling implication of the consent-requirement if it prohibited imposing minor defensive harms on lethal aggressors, I take it to be a virtue of my account that it gives a principled rationale for why there might plausibly be an exception in these cases.⁶⁹

4.3 Explaining the verdict in *Elevator 2*

The PPE also neatly accounts for the intuitive judgement in cases like *Elevator 2*. In assessing this case, I propose that we proceed by composition, treating each of the five victims in the same way that we treat the single victim in *Elevator 1*. This is in line with the broadly liberal/individualist moral outlook endorsed by consent sceptics. The basic

⁶⁶ I should acknowledge that an anonymous associate editor reports having a less clear intuition about the wrongness of intervention in *Poisoned Patient*.

⁶⁷ Thanks to an anonymous associate editor for pressing me to consider these cases.

⁶⁸ For a detailed discussion of how aggressor-centric reasons may contribute to the justification of defensive harm, see Kimberly Brownlee, 'Acting Defensively for the Sake of Our Attacker' (unpublished manuscript).

⁶⁹ Interestingly, the rationale for this exception may extend to certain cases of killing, depending on our views about how these justifications for harm grounded in preventing injustice aggregate. For example, imagine a case in which an aggressor threatens to kill 100 innocent victims, who each validly refuse defensive intervention. If we grant that a rescuer would be permitted to inflict a minor defensive harm on the aggressor in order to prevent the injustice of one murder, it is possible that the amount of harm it is permissible to inflict increases with each additional threatened murder, to the point at which it is permissible to kill the aggressor. Though I am open to this idea, its plausibility depends on resolving some extremely difficult issues concerning how harms and injustices aggregate, which I lack the space (and skill) to explore here. Thanks to an anonymous associate editor for raising this possibility.

thought is that the normative consequences that follow from each victim's consent and refusal should not change simply because they are threatened in conjunction with others.⁷⁰ On this approach, each of the four refusing victims successfully exercises their normative power to make their good morally inadmissible for the purposes of justifying defence, just as the single victim does in *Elevator 1*. All that changes between *Elevator 1* and *Elevator 2* is whether this exercise renders defence impermissible. It does not, due to an important feature of *Elevator 2*: each victim's good is *sufficient* to justify the total defensive harm involved in defending the group (killing Aggressor). This explains the intuitive permissibility of defending the five victims, despite the fact that four refuse. Since the good of each victim is sufficient to justify the total defensive harm, and one victim consents, Rescuer has sufficient admissible moral reason to justify killing Aggressor.

This account also shows that defending the five in *Elevator 2* need not involve paternalistically wronging any victim. Though her action is not endorsed by four of the victims, her justification for defending them does not require making use of any victim's moral resources without their permission (unlike in *Elevator 1*). It is not true of any victim that they are defended against their will, for the sake of *their own* good. Rescuer thus respects each victim's sovereignty over their good.⁷¹

4.4 Consent and Proportionality

With the foregoing analysis of the two cases to hand, we are in a position to see why, *contra* the sceptics, the consent requirement imposes a significant constraint on defending groups. The key point is that there is a crucial difference between cases like *Elevator 1* and *Elevator 2*, and multiple-victim cases more generally. This is that in many cases (including humanitarian interventions, rebellions, and wars of national-defence) defending victims not only involves harming aggressors (as in *Elevator 1* and *Elevator 2*) but also causing serious harm to innocents.

This is significant because, when there are these additional harms, it is *not* true (as in *Elevator 1* and *Elevator 2*) that each individual victim's good is sufficient to justify the

⁷⁰ I will consider an important challenge to this assumption in Section 5.

⁷¹ Allen Buchanan has recently argued that we should reject the consent requirement in multiple-victim cases, on the ground that basing the requirement in anti-paternalism would, implausibly, require unanimous consent. My account shows that this objection is mistaken. One does not treat others paternalistically simply by acting against their will, but by justifying one's action by appeal to certain reasons: those pertaining to their good or welfare. Since, as I have argued, it is possible to justify defending non-consenting victims without appealing to their good, anti-paternalism does not require unanimous consent (or non-refusal). See Buchanan, 'Self-Determination, Revolution, and Intervention', n.10. and 'The Ethics of Revolution and Its Implications for the Ethics of Intervention', p.317.

total amount of harm that would be caused by defending them. Instead, when defence involves collateral killing the *number* of victims becomes highly morally relevant, due to considerations of proportionality.⁷² Putting aside some important complications, the proportionality requirement holds that the relevant goods achieved by a defensive action must be sufficient to outweigh the relevant bad effects that it causes.⁷³ When those bad effects include killing innocent persons, defence is usually permissible only if it will prevent a significantly greater number of innocent people from being killed. In these cases, the rights of those collaterally killed are justifiably overridden as the ‘lesser-evil’.⁷⁴ For example, judgements about the standard Trolley Case suggest that a 5:1 ratio between lives saved and innocents killed might be required. At the very least, it would not be proportionate to kill one innocent person as a side-effect in order to save one innocent person from being killed.

Given the proportionality requirement on harming innocents, defensive actions that cause collateral harms are often impermissible, because disproportionate, in virtue of a *lack of victims*. My contention is that a lack of *consenting* victims (or, more accurately, a lack of *non-refusing* victims) can also render defence disproportionate, for precisely the same reason: defence would not achieve sufficient relevant goods to outweigh its relevant bad effects. Just as in *Elevator 1* and *Elevator 2* discussed above, the central idea is that when victims validly refuse defensive intervention, they thereby exclude their good from the pool of reasons that rescuers may appeal to in order to justify defending them. However, in cases where, unlike *Elevator 1* and *Elevator 2*, the number of victims matters because defence involves killing innocents, the inadmissibility of these reasons may render defensive force disproportionate, depending on the specific facts of the case. If correct, this reveals the sceptical aggregation principle – the Minimal Consent Requirement – to be mistaken. There are many cases in which victims’ refusal can make it impermissible to defend a group, even if some (or even many) members explicitly consent to defensive intervention.

To illustrate this, consider a variation on *Elevator 2* that involves collateral killing:

⁷² More precisely, I am concerned here with what Jeff McMahan calls the ‘wide’ proportionality constraint, which governs harms inflicted on non-liable persons. Harms to liable are governed by a ‘narrow’ proportionality constraint persons (at least up to certain point). See McMahan, *Killing in War*, pp.18-20.

⁷³ For a classic discussion, see Thomas Hurka, ‘Proportionality in the Morality of War’, *Philosophy and Public Affairs* 33, No.1 (2005), 34-66.

⁷⁴ As Helen Frowe explains: ‘Lesser-evil justifications obtain when one will prevent substantially more harm than one causes, such that the disparity between the harm and the good overrides the deontological presumption against causing harm’. Helen Frowe, ‘Claim Rights, Duties and Lesser-Evil Justifications’, *Proceedings of the Aristotelian Society* 89, No.1 (2015), 267–85, at p.274.

Elevator 3: Five Victims are taking a ride in an elevator. Aggressor begins to saw through the elevator cable because he hates Victims and wishes them dead. Rescuer is walking by and is able to kill Aggressor by throwing a grenade, thereby saving Victims. However, if she does so, the blast from the grenade will kill Bystander as a side-effect. While one Victim consents to Rescuer's intervention, the four remaining Victims refuse because of their commitment to non-violence.

In this case, I contend, it is impermissible for Rescuer to defend the five Victims, because doing so would be disproportionate. This is despite the fact that saving the five *would* be permissible, because proportionate, if each of the Victims consented (or did not refuse).⁷⁵ In essence, on my account, the four Victims' refusal morally converts a case of collaterally killing one in order to save five into a case of collaterally killing one in order to save one. And saving a single life is not enough to justify killing one innocent person.

Three features of my proposal are worth highlighting. First, on this view consent functions in the morality of defensive harm as an internal component of the broader proportionality requirement on harming innocent people. An individual's consent (or non-refusal) is required in order for benefits to that individual to count positively in the proportionality calculation. Second, the proposal is emphatically not majoritarian. What matters morally is the *number* of refusing victims, not the *proportion* of refusers within a victim group. Third, on this view, what counts as having sufficient consent to satisfy the consent requirement will vary on a case-by-case basis. This is because the amount of consent required in order to render defence proportionate will depend on contingent factors, chiefly (i) how large the group of beneficiaries is, (ii) how many victims validly refuse defensive intervention, and (iii) the extent to which innocents will be harmed in the course of defending the group. When the number of victims far exceeds the threshold at which causing a certain amount of harm to innocents would be proportionate, only a small proportion of victims need consent (or not validly refuse) in order to render it proportionate to cause that harm. But as the number is reduced, getting closer to that threshold, the proportion of victims that must consent correspondingly increases, potentially to the point at which a *single* refusal would render defence impermissible (this may be true of *Elevator 3*, for example). My view is thus both weaker and stronger than the Majority Consent Requirement.

⁷⁵ I am assuming here that saving five innocent persons is the threshold at which collaterally killing one innocent person becomes proportionate. But these numbers are just for illustration. One may substitute whichever threshold ratio one wishes.

The account that I have proposed can be captured more precisely in the following aggregation principle:

The Proportionate Consent Requirement: Other-defence is morally justified only if a sufficient sub-set of the members of the victim group consent to defensive intervention (or do not validly refuse). A sub-set is sufficient if the benefits of defence to the members of the sub-set are sufficient to justify the total amount of morally weighted harm that defending the group will cause.

Here is a slightly more intuitive way of understanding this interpretation of the requirement. Since, on my account, a victim's refusal serves to render the fact that they would benefit inadmissible for the purposes of establishing lesser-evil justifications for harm innocents, there is an important sense in which we can treat victims' refusal as morally equivalent to their non-existence. Given this, one way of determining whether the consent requirement has been met in a particular case is by applying the following counterfactual test: If the (validly) refusing member(s) of a victim group did not exist, would it be permissible to cause the same harm to innocents in order to defend the remaining members? If the answer is no, the consent requirement has not been met and defence is impermissible.

5. Obligations to Consent? Wrongful Refusal?

A central idea behind the Proportionate Consent Requirement is that we should take each member of a victim group in isolation, so that the normative consequences that follow from each victim's refusal are unaffected by the presence of the other victims. Put differently, we should treat each victim in a multiple-victim case in the same way that we would treat a single victim, and then proceed additively. Given this, one important line of objection to my proposal holds that there is some factor present in multiple-victim cases, absent in single-victim cases, which undermines this moral independence.

Here is one obvious candidate: in multiple-victim cases each victim's decision to refuse defensive intervention may significantly affect *other victims'* prospects. This is especially salient on my account. Consider the following variation on *Elevator 3*:

Elevator 3*: Five Victims are taking a ride in an elevator. Aggressor begins to saw through the elevator cable because he hates Victims and wishes them dead.

Rescuer is walking by and is able to kill Aggressor by throwing a grenade, thereby saving Victims. However, if she does so, the blast from the grenade will kill Bystander as a side-effect. Four Victims consent to Rescuer's intervention, while one Victim (let's call him Rafi) refuses because of his commitment to non-violence.

According to the Proportionate Consent Requirement, Rafi's refusal makes it impermissible for Rescuer to defend the group. By exercising his PPE, Rafi reduces the goods available to Rescuer to below the threshold at which she would have a lesser-evil justification for killing the bystander. Rafi's decision is thus a matter of life or death for the remaining four Victims (on the assumption that Rescuer will only defend them if she is morally justified in doing so). This other-regarding feature of multiple-victim cases plausibly undermines the idea that each victim's consent and refusal are morally independent. Instead, we might object that the fact that Rafi's refusal would have serious negative consequences for others places limits on its moral effectiveness.

More precisely, the objection under consideration consists in two distinct claims. The first is that Rafi would act morally wrongly by refusing defensive intervention. In fact, he is *morally required* to consent, because doing so will bring about the saving of four consenting victims (by furnishing Rescuer with a justification for defending them) and – by hypothesis – we have duties to perform easy rescues.⁷⁶ The second claim holds, adopting an idea of David Estlund's, that the fact that it would be wrong for Rafi to refuse renders his refusal morally invalid. In virtue of this invalidity, the normative situation is as it would have been had Rafi in fact consented.⁷⁷ For our purposes, what this means is that Rafi's good *is* in fact morally admissible for justifying defensive harm, because his attempt to exercise his PPE is unsuccessful. Since Rafi is morally required to make his good available for the sake of his co-victims, Rescuer may use Rafi's moral resources without his authorisation. Hence – *contra* my proposal – Rescuer *does* in fact possess a lesser-evil justification for defending the group.

This objection is notable in that it aims to show that the independence assumption is mistaken from within a broadly liberal/individualist picture of morality. It does not claim

⁷⁶ A roughly parallel line of objection is suggested by Victor Tadros, who argues that victims of crime may have duties to authorise the punishment of offenders in order to protect future potential victims of crime. Tadros, *The Ends of Harm*, pp.297-298.

⁷⁷ David Estlund, *Democratic Authority* (Princeton: Princeton University Press, 2008), Ch.8. As Gopal Sreenivasan puts it, this view “treats the fact that Y *should have* consented to X's ϕ -ing as entailing the normative upshot that would have resulted had Y actually consented.” Gopal Sreenivasan, “Oh, but you should have”: Estlund on Normative Consent’, *Iyyun: The Jerusalem Philosophical Quarterly* 58 (2009), 62-72, at p.67.

that there is anything morally special about defending groups, over and above the individuals that compose them. Instead, it holds that the moral independence of victims is undermined simply by individuals' obligations to attend to the welfare of others. If successful, it is a mistake to claim that we may only justify defending groups of victims by appeal to benefits to victims who have not refused defensive intervention, because victims' normative power to refuse is importantly limited by the effects that its exercise would have on others. In fact, if pushed far enough, the objection could be employed in defence of the sceptics' Minimal Consent Requirement. On this view, if *any* member of a victim group consents to being defended (or does not validly refuse), *all* other members are required to consent in order to bring this about, and so their refusal is morally invalid.

Though intuitively powerful, I believe this objection fails. Its force comes from an analogy between what I have called our 'moral' resources – the moral reasons arising from one's good or welfare – and ordinary material resources over which we have powers of exclusion, such as our bodies and property. Given this, one straightforward response would be to accept the analogy, but deny that individuals can be subject to duties to use their resources – both material and moral – to assist others. But this is a rather extreme and unpalatable position. When it comes to material resources, it is very intuitive that individuals can be morally required to make life-saving contributions. It also seems plausible that, in at least some cases, refusing to contribute resources would be morally ineffective, and so others may use those resources without wronging their owner. For example, if I can easily save four lives by using my body to pull four non-swimmers from a shallow pond, I may be duty-bound to do so and others may use my body for that purpose if I refuse to do so myself without violating any claim of mine.⁷⁸

A more plausible line of response is to deny the resource analogy. We may accept that we are subject to rescue duties with respect to our material resources, but deny that we are similarly obligated when it comes to our moral resources. One way to argue for this is to identify differences in the burdensomeness of the two kinds of duty, since duties to assist are typically constrained by the costs to the duty-holder of discharging them. For example, we might hold that requiring a committed pacifist to use his own good in order to bring about killing would involve him acting against his conscience in a particularly egregious way.⁷⁹ By contrast, a duty to use one's property to pull non-swimmers from shallow ponds would not. An obvious difficulty for this response, however, is that it

⁷⁸ For detailed discussion of the connection between duties to aid and the permissibility of using others as a means, see Tadros, *The Ends of Harm*.

⁷⁹ Thanks to anonymous associate editor for suggesting this response.

requires a plausible explanation of why the importance of preserving one's moral integrity should trump others' interest in being protected from murder. There is something troubling about the idea that claims of conscience can enable us to evade duties to others.⁸⁰

I too think the objection fails due to a disanalogy between material and moral resources, but for a different, and more fundamental, reason. The previous response grants the assumption that victims have positive moral reasons to use their good for the sake of their co-victims in cases like *Elevator 3**, but argues that these reasons do not generate a duty. By contrast, I argue that victims do not have *any* moral reason to use their good for the sake of others in cases like this. This is because we have no moral reason to create lesser-evil justifications for harming innocent people. If there is no reason to rescue others *in this particular way*, then victims cannot be under a duty to consent to defensive intervention, nor act wrongly by refusing to do so. Hence, there are no grounds for invalidating their refusal.⁸¹

To help see this, it is useful to think about other ways in which we can use our good to create lesser-evil justifications, other than by consenting to third-party intervention. Consider the following case:

Trolley: Aggressor has set a trolley rolling towards four innocent victims who are tied to the main tracks, because she hates them and wishes them dead. Rescuer can divert the trolley onto a side-track, thereby saving the four. However, there is an innocent bystander on the side-track who will be killed if she does so. Further down the main track, Barry is crossing the tracks on his way to work and sees events unfolding up the line. While he is crossing, the trolley also threatens his life, but he can avoid being threatened by stepping off the tracks.

Assume, once more, that it is permissible to collaterally kill one innocent person in order to save five lives, but not to save four, and that Rescuer will only divert the trolley if she is justified in doing so. Barry then has two options. First, he can remain on the tracks and thereby bring about the saving of the four (by making his good available to justify Rescuer diverting the trolley). Second, he can step off the tracks and thereby fail to bring about the saving of the four (making his good unavailable to Rescuer.) Would Barry act wrongly by

⁸⁰ For a detailed version of this objection as applied to duties to defend others from unjust attack, see Fabre, 'Mandatory Rescue Killings', pp.378-384.

⁸¹ This is not to deny that refusal may be morally wrong, and therefore potentially invalid, on different other-regarding grounds. For example, if a victim refuses defensive intervention purely out of racist hatred for their co-victims. My claim is that refusal is not invalidated *simply* by the fact that its validity would result in other victims not being defended. For discussion of various other-regarding constraints on consent, see Grill, 'Liberalism, Altruism, and Group Consent.'

stepping off the tracks? Does he have any moral reason to stay put? It seems clear (at least to me) that Barry has no such reason (for reasons I will elaborate on below). To emphasise this, imagine that Barry has already crossed the tracks before becoming aware of the situation up the line. Does he have any reason to step back onto the tracks in order to bring about the saving of the four? It seems absurd to think that he does.

This case is intended to be relevantly analogous to *Elevator 3**, in which four victims consent to defence while one (Rafi) refuses. Both Rafi and Barry have the ability to save four lives at the cost of one life by contributing their moral resources (by consenting to defensive intervention and by staying put, respectively). But if Barry has no reason to do so, this suggests the same should be true of Rafi. If that's right, then, *contra* the objection under consideration, victims do not act wrongly by refusing defensive intervention, even if this results in others not being defended from serious harm. Their refusal is therefore not invalidated and their good remains morally unavailable to rescuers, in line with The Proportionate Consent Requirement.

One might object that there is an important difference between cases like *Elevator 3** and *Trolley*.⁸² In the former, Rafi has no control over whether he faces a threat, whereas in the latter Barry has a choice. This could be significant in two ways. First, individuals are not usually required to aid others if doing so would expose them to a threat of serious harm. If that is right, then we have a straightforward explanation of why Barry is under no obligation to create a lesser-evil justification that does not apply in Rafi's case.

The problem with this objection is that even if it shows that Barry is *not required* to step onto the tracks, it does nothing to explain the intuition that he *lacks a moral reason* to do so. Even if Barry were perfectly willing to expose himself to the threat in order to save the four, it's still a bizarre thing to do. The objection therefore does not challenge my claim that victims like Rafi similarly have no moral reason to create lesser-evil justifications by consenting to defensive intervention.

Second, and more promisingly, we might hold that an individual's claim to defensive intervention is vitiated if they intentionally choose to expose themselves to a threat. If that is true, then Barry's attempting to create a lesser-evil justification for saving the four by stepping onto the tracks would be self-defeating, since it would fail to add an additional claim to be rescued to those of the four. One might object that *this* is the explanation of why Barry not only lacks a duty but also a moral reason to create a justification for saving the four, which does not apply to *involuntary* victims such as Rafi.

⁸² Thanks to an anonymous associate editor for pressing me to consider the following two objections.

Hence, the fact that Barry clearly lacks the relevant reasons does not show that victims of aggression also lack them.

However, we can see that considerations of self-defeat fail to explain Barry's lack of reasons by considering cases in which he can trigger a lesser-evil justification without exposing himself to a threat:

Trolley 2: Aggressor has set a trolley rolling towards four innocent victims who are tied to the main tracks, because she hates them and wishes them dead. Rescuer can divert the trolley onto a side-track, thereby saving the four. However, there is an innocent bystander on the side-track who will be killed if she does so. Barry is passing the tracks on his way to work and sees events unfolding up the line. Jimmy is sound asleep on a nearby bench. Barry is able to gently place Jimmy on the main track without waking him.

Again, Barry clearly has no moral reason to manufacture a justification for saving the four at the cost of the bystander's life. But this is not because it would be self-defeating to try. Since Jimmy would not have voluntarily placed himself in harm's way his claim to be rescued would be undiminished, and so Rescuer *would* acquire a justification for diverting the trolley as a result of Barry's action.

One might respond that it would be independently morally wrong for Barry to use Jimmy in this way, and that this explains Barry's lacking reason to perform it so. But this seems to misrepresent the normative situation. In normal cases of wrongfully using someone as an instrument to promote the good, the agent has a perfectly intelligible reason to do so, which *is then defeated* by the constraint on using others as a means. This doesn't seem true of Barry; he has no such reason in the first place. Using Jimmy wouldn't simply be wrong, it wouldn't *make sense*.

Putting aside intuitions about particular examples, here is a more general diagnosis of what makes these cases so odd. The objection that victims can be morally obligated to consent, in order to generate justifications for harming the innocent, relies on the idea that we can have moral reasons to alter the balance of moral reasons. This is not always problematic. For example, if Barry could push the innocent bystander off the side-track, thereby removing the moral barrier preventing Rescuer from saving the four, then he seems clearly justified in doing so. But when it comes to altering the balance of reasons *in order to bring it about that rights against harm are overridden* something has gone seriously wrong. For if the pre-existing reasons in favour of killing the innocent are insufficiently weighty to justify doing so, how can this generate any reason to *add* additional considerations in

favour of harming to the scales, in order to tip the balance?⁸³ To claim that there is such a reason requires counting the reasons in favour twice. That's not only irrational, but also involves an objectionable attempt to circumvent the moral constraints that are properly imposed by individuals' basic rights. It is this general failure of practical reasoning, I contend, that explains why victims have no reason to make their good available to rescuers.

Let me elaborate in the context of *Elevator 3**. The idea that Rafi is morally obligated to contribute his good in order to rescue his four co-victims seems compelling when we focus on the claims of the four who are in dire need. But this obscures the crucial fact that the weight of those claims has been fully accounted for in the moral calculus, and found to be *defeated* by the bystander's claim not to be killed. We therefore cannot – on pain of double-counting – appeal to those claims in order to ground Rafi's obligation to consent, and there are no remaining other-regarding reasons that could do the job. If anything, the claim of the bystander gives Rafi a reason *not* to consent.

In summary: While we may have duties to help others perform independently justified actions – by using our bodies and property, for example – we do not have reason to do so by manufacturing lesser-evil justifications for their actions. Since the proposed objection to the Proportionate Consent Requirement relies on the claim that we do have such reasons, it should be rejected.

6. Conclusion: The Ethics of Protest

By way of conclusion, let me sketch how my account of the role of consent in the morality of defensive harm may illuminate some wider practical issues, by looking at the case of political protest.

Acts of public protest can have a variety of aims. In many cases the main goal seems to be *indicative*, in that the protesters are trying to draw attention to considerations that they believe decisively count against an action or policy. They can be interpreted as (loudly) giving others a report on their interpretation of the facts. But this is not the only function of protest. As well as providing information about the world, protest also seems to have a distinct *expressive* aim, in that the protesters are trying to communicate an attitude

⁸³ Here is another way of bringing out this contrast, which focuses on our reasons to have certain attitudes, rather than our reasons to act: it seems perfectly reasonable for an uninvolved third-party to *regret* the presence of the bystander on the side-track in *Trolley*, who prevents the four from being saved. But it seems very odd to regret the absence of an additional victim on the main-track, whose presence would make it permissible to save the four by killing the bystander.

of personal disapproval towards the object of protest. In addition, the account of consent and refusal that I have sketched suggests that protest may also have a third, *transformative* function, whereby the act of protest *itself* counts against the justification of the action or policy being protested. Protesters may be interpreted not simply as giving a report or expressing an attitude, but also as exercising a normative power: the power to withhold their consent to the policy being protested, and thereby prevent the fact that the policy may benefit them from contributing to its justification.

The protests in Britain against the 2003 invasion of Iraq (or at least a slightly idealised version) are a useful case to illustrate this idea.⁸⁴ The protesters' rallying cry of "Not in My Name!" can be understood in the indicative, expressive, and transformative senses outlined above.⁸⁵ Clearly, one thing the protesters were trying to do was to highlight certain perceived facts about the war – that it lacked a just cause, for example. But the protest also had a self-regarding aspect to it; the protesters were communicating that they did not endorse the war being fought on their behalf. If we treat this as merely expressive, it may seem morally trivial, or even self-indulgent ("It's not all about *you!*"). However, if we understand protest as having a transformative function, as I have suggested, it may have considerable moral significance. One of the main justifications given (at least initially) by the British government for participating in the invasion was that doing so would help protect Britain and its residents from attack by WMDs and terrorist groups harboured by the Iraqi regime.⁸⁶ If we interpret the anti-war protesters as exercising a power to withhold consent, one consequence of their protest was to prevent a portion of that putative benefit from contributing to the justification of the invasion. On this view, the existence of protest may *itself* have made the invasion harder to justify than it otherwise would have been, independently of the other factors that may have counted against it.

⁸⁴ For an illuminating discussion of this case, and of anti-war protest in general, see John Williams, "Not in My Name": Legitimate Authority and Liberal Just War Theory' in Anthony F. Lang Jr., Cian O'Driscoll, and John Williams (eds), *Just War: Authority, Tradition, and Practice* (Washington DC: Georgetown University Press, 2013), 63-80. Williams also understands the moral significance of protest in terms of withdrawing consent, though his account of the normative consequences of withdrawal differs from my own.

⁸⁵ For a different take on the relevance of the protesters' slogan to the permissibility of force, see Pattison, 'Representativeness and Humanitarian Intervention', p.577.

⁸⁶ As Tony Blair made clear in his address to the nation on the eve of the invasion: "The threat to Britain today is not that of my father's generation. War between the big powers is unlikely, Europe is at peace, the Cold War already a memory. But this new world faces a new threat of disorder and chaos born either of brutal states like Iraq armed with weapons of mass destruction or of extreme terrorist groups. Both hate our way of life, our freedom, our democracy. My fear, deeply held, based in part on the intelligence that I see is that these threats come together and deliver catastrophe to our country and our world." Prime Minister's Address to the Nation. March 20th 2003. Full transcript available at http://news.bbc.co.uk/1/hi/uk_politics/2870581.stm [last accessed 20th December 2016].

This last point highlights that the transformative function of protest can be morally successful even if the indicative and expressive aspects fail (for example, because the action or policy lacks independent moral defects for protesters to highlight, or because the protestors are entirely ignored by their intended audience). Even in these cases, the mere fact of protest may successfully diminish the justification for the action or policy. In some cases, the existence of protest may even be *sufficient* to render an action or policy morally unjustified, independently of its other moral merits or defects.