HHS Public Access

Author manuscript

Bioethics. Author manuscript; available in PMC 2021 November 01.

Published in final edited form as:

Bioethics. 2020 November ; 34(9): 906–911. doi:10.1111/bioe.12673.

Bystanders, Risks, and Consent

Helen Frowe

Stockholm University, Dept. of Philosophy, Stockholm, Sweden

Abstract

This paper considers the moral status of bystanders affected by medical research trials. Recent proposals advocate a very low threshold of permissible risk imposition upon bystanders that is insensitive to the prospective benefits of the trial, in part because we typically lack bystanders' consent. I argue that the correct threshold of permissible risk will be sensitive to the prospective gains of the trial. I further argue that one does not always need a person's consent to expose her to significant risks of even serious harm for the sake of others. That we typically need the consent of participants is explained by the fact that trials risk harmfully using participants. Bystanders, in constrast, are harmed as a side-effect, which is easier to justify. I then consider whether the degree of risk that a trial may impose on a bystander is sensitive to whether she is a prospective beneficiary of that trial.

Keywords

Suggestions: Research ethics; Human subjects research; Research subjects; Informed consent; Risk; Bystanders

INTRODUCTION

Remarkably little attention has been paid to the moral status of *affected bystanders* in medical research trials[1] – that is, the status of people who are not enrolled in trials, but who are exposed to risks of harm by those trials.[2] The comparative lack of literature (and legal protection) is especially surprising given the extensive writings on the protections afforded to participants in such trials.[3] Participants, in contrast, must receive comprehensive information about the trial, must explicitly consent to participate, typically have some degree of control over how information about them is used, have stringent privacy rights regarding their participation in the trial, may withdraw from the trial at any time, and are often financially compensated. These protections are understandable, since medical trials expose participants to risks of harm. But trials also expose bystanders to risks of harm.

This paper argues that we should employ a *ratio* account of permissible risk imposition for bystanders in clinical trials.[4] Ratio accounts weigh the prospective risks against the prospective benefits of the trial. They contrast with *threshold* accounts, which set a universal limit to risk imposition irrespective of the prospective gains of any particular trial. Two

Corresponding Author: Helen Frowe, University of Stockholm, Dept. of Philosophy, Huset D, Frescati Campus, Stockholm SE-10691, Sweden, helen.frowe@philosophy.su.se.

recently proposed threshold accounts also advocate very low limits on risk imposition, arguing it is impermissible to impose any significant risk on bystanders at all, at least in part because we typically lack their consent.[5] I argue that these proposed limits are too restrictive, and that one does not always require a person's consent to expose her to even serious risks and serious harms.

The paper proceeds as follows. Section Two outlines two threshold accounts: the Ordinary Life Standard and the Near Zero Standard. Section Three argues that threshold accounts are implausible when considered in a broader context of permissible risk imposition, and that a more plausible account will rely on a comparison of harm to expected benefit.[6] I further argue that comparison accounts will sometimes allow even significant risk imposition. This gives us further grounds to reject Ordinary Life Standard and the Near Zero Standard, which permit only low risk impositions, regardless of circumstance. Section Four argues that we require the consent of trial participants not because they are exposed to risks of harm, but because they are exposed to a risk of being harmfully used. Bystanders, in contrast, are not harmed used, but harmed as a side-effect. Section Five considers the role of consent in making more significant collateral harms permissible. Section Six considers how the degree of permissible risk imposition might be further influenced by whether a person is an expected beneficiary of the risk-imposing activity. Section Seven concludes.

THRESHOLD APPROACHES TO PERMISSIBLE RISK

The Ordinary Life Standard

Battin et al. offer perhaps the most developed account of risks to bystanders in medical trials. They defend a universal threshold for permissible risk imposition on identifiable persons, arguing that if a trial poses a risk of non-trivial harm to an identifiable person, then the trial requires her consent if she cannot easily avoid the risk.

The notion of non-trivial risk at work in this account is most plausibly interpreted as a function of two things. The first is the *degree of possible harm* to the bystander: trials that might require consent are those that expose bystanders to potentially serious harm, such as "contracting a disease with significant morbidity or mortality".[7] The second is the *degree of increased risk*: one needs consent if the risk that a bystander will suffer the potentially serious harm exceeds her risk of suffering that harm in ordinary life. I'll refer to the view that consent is required when a trial imposes a risk of a harm that exceeds the risk of that harm in ordinary life as the *Ordinary Life Standard* for consent.

By way of illustration, consider that in the United States a person has roughly a .000015 chance of contracting malaria. If we take that as a benchmark for the risks of malaria in ordinary life, exposing a bystander in the US to a .00002 chance of contracting malaria as a side-effect of a trial would be to impose an increased, non-trivial risk upon her. By Battin et al.'s lights, we must obtain the consent of bystanders who are exposed to such risks before proceeding with the trial.

The Near-Zero Standard

In February 2017, a report commissioned by the National Institute of Allergy and Infectious Diseases and the Walter Reed Army Institute of Research in the United States argued that a human challenge trial on the Zika virus would not currently be ethical. The Zika virus, which is transmitted through sexual contact, mosquito bites, and other transfer of bodily fluids, causes severe birth defects. The World Health Organisation declared an outbreak of Zika in 2016 a global public health emergency. The report's authors rejected the feasibility of an ethical trial in part because it would expose bystanders to risk of infection. And, they argued, "because these third parties generally cannot know about, protect themselves from, or consent to risks, risks to them are only reasonable if they can be reduced to near-zero."[8] Call this view – that virtually no risks to bystanders are justified without their knowledge or consent – the *Near Zero Standard*.

Both the Ordinary Life Standard and the Near Zero Standard are *threshold accounts*: that is, they set a standard of permissible risk imposition that is insensitive to the prospective benefits of the trial, focusing only on the risk to bystanders. Further, both treat consent (or the lack thereof) as at least partly justifying the very low threshold of permissible risk imposition.

AGAINST THRESHOLD ACCOUNTS

Consider *Leg*:

> *Leg*: A runaway trolley is heading to where it will very likely kill one person. Onlooker can divert the trolley down a side-track, where it is equally likely to cause Workman to lose his leg.

Say that we agree that it is impermissible to impose the high risk of the loss of a leg on Workman in the course of saving one life – that is, that this cost exceeds what Workman can be required to bear for the sake of the one. That might be plausible. What is implausible is that the cost Workman can be required to bear is wholly insensitive to how many people will be killed by the trolley – that even if twenty people will be killed, Onlooker may not divert. And yet this is what threshold accounts propose: once we exceed the threshold, the amount of harm at stake for others simply makes no difference to whether Workman may be made to suffer the loss of his leg. This is particularly worrying with respect to the Ordinary Life Standard, since on that view the converse also holds. If one requires consent when one increases the risk of only serious harms, one may inflict *non*-serious harms without consent no matter how trivial the prospective benefit of the trial.

It is instructive that we do not use threshold standards to assess risks to participants in trials. Rather, risks to participants are justified only if the trial satisfies a cost-benefit analysis – that is, only if the risk to the participants is outweighed by the prospective benefit of that particular trial.[9] To simplify, I'll call this a *ratio* standard for risk imposition (although the correct form of comparison may be more complex than a straightforward ratio, of course). Given the use of a ratio standard for participant risk, why would we adopt a threshold standard for bystander risk?

The adopting of this radically different approach may be underpinned by an assumption that one may risk harm on the basis of a cost-benefit analysis only when the potential victim has consented to be exposed to risk. Without consent – which we typically lack from bystanders – one may risk only very minimal harms, irrespective of the good that one might achieve. But this seems very *ad hoc*: we do not apply this standard in other risk-imposing contexts in which we lack consent. For example, consider *Trolley*:

> *Trolley*: A runaway trolley is heading to where it will kill five people. Onlooker can divert the trolley down a side-track, where it will kill Workman.[10]

Most people think that it is permissible to divert the trolley towards Workman (and even those who are uncertain about the permissibility of diverting when five lives are at stake are likely to grant that diverting would be permissible for the sake of a higher number of lives – perhaps ten or twenty).[11] But this permissibility does not assume that Workman has consented to be killed – on the contrary, we typically assume that he has *not* consented.

Of course, we might reply that this approach to defensive harm is mistakenly permissive. Perhaps we should not expose people to such high risks of harm for the sake of others. But reflecting on other risky activities also supports the more permissive standard. An ambulance driver or police officer, for example, may drive more quickly than others when they are rushing someone to the hospital or chasing a dangerous suspect. Even though driving faster increases the chances of killing a bystander, this increased risk can be permissible when weighed against the good of saving a life or catching a criminal who threatens harm to others.

There are various possible explanations of why we permit such risky activities. We might, for example, point to a type of social contract (e.g. that we all want to have access to ambulances, and so accept the increased risk from their use for others). But these explanations could also apply to clinical trials. We all benefit from advances in medicine, either directly or indirectly. Many people would want to participate in a trial if doing so might save their life.

If engaging in these risk-imposing activities is permissible, it follows that one may risk very serious harm to a person, without her consent, if there is sufficient good at stake. The plausibility of threshold accounts is thus significantly undermined. What we need is a standard that weighs prospective harms against prospective goods, as we see in the proportionality calculations in cases of defensive harm.

These comparisons with other risk-imposing activities further suggest that the particular thresholds employed by the Ordinary Life Standard and the Near Zero Standard are far too restrictive. Workman's chances of being killed in ordinary life are irrelevant to the proportionality of killing him in *Trolley* (even if we could work out what counts as 'ordinary life', which seems a rather difficult task).[12] Proportionality (roughly) compares the inflicted harm with the averted harm.[13] It is not at all sensitive to the risks of ordinary life. It is not, for example, proportionate to kill a person for the sake of saving two lives if she lives in a comparatively dangerous country, but disproportionate to do so if she lives in a

comparatively safe country. And it certainly does not limit the permissible risk of non-consensual harming to near zero.

USING AS A MEANS

That one may impose some non-consensual harms on a person without her consent does not mean that consent is always irrelevant to permissible risk imposition. It is often relevant. For example, one typically requires the consent of a participant in a trial. But this requirement is not grounded in the mere fact that a participant risks being harmed, but rather in the fact that she risks being harmfully used. Harmfully using a person is much harder to justify than harming her as a side-effect. Consent is therefore much more often necessary for justifying harmfully using a person than for harming her as a side-effect. Contrast *Trolley* with *Push*:

> *Push*: A runaway trolley is heading to where it will kill five people. Onlooker can lethally push Hiker in front of the trolley as it passes, so that his body blocks the trolley and saves the five.

Most people think that it is impermissible to push Hiker onto the tracks, using his body as a human shield for the sake of the five.[14] Hiker is analogous to a participant in a trial: one does not aim at harming him *per se* (one would be pleased if he survived the trolley), but one makes *use* of him in a way that (potentially) harms him. Since harmfully using is very hard to justify, it is typically permissible only if the person consents. But collateral harms are less egregious, and may more often be imposed without consent.

NARROW AND WIDE PROPORTIONALITY

A fairly standard view of permissible harm to bystanders is that the harm is justified only if it is the lesser evil – that is, only if it is significantly outweighed by the good we can secure, taking into account various deontological features, such as the fact that it is a harm that is caused rather than allowed, whether it is merely foreseen rather than intended, and so on.[15] And, as just noted, it is typically thought that I don't need the consent of the bystander in a case of proportionate harming.

But if the harm to the bystander were *dis*proportionate, and as such did, *prima facie*, require her consent, it is unclear why we would want to inflict it. After all, if 'disproportionate' means that it is not outweighed by the good we hope to secure, it looks irrational to inflict it. We can, however, draw on Jeff McMahan's work on proportionality to resolve this puzzle. Consider *Attack* and *Collateral*:

> *Attack*: Murderer is trying to unjustly kill Victim, who is innocent. Victim can prevent Murderer's attack by throwing a grenade at Murderer. The blast will kill Murderer, and temporarily knock out Bystander (but do Bystander no lasting harm).
>
> *Collateral*: Murderer is trying to unjustly kill Victim, who is innocent. Victim can prevent Murderer's attack by throwing a grenade at Murderer. The blast will kill Bystander, and temporarily knock out Murderer (but do Murderer no lasting harm).

It seems clear enough that whilst it would be proportionate for Victim to kill Murderer, it is disproportionate for him to kill Bystander. And yet, in both cases, Victim is fending off a

threat to his life, with the same amount of harm (killing one and rendering one unconscious). This suggests that proportionality is sensitive to the moral status of who suffers the harm. As a culpable attacker, Murderer forfeits his right not to suffer harm proportionate to the threat that he poses – that is, he is *liable* to harm proportionate to the threat he poses.[16]

McMahan suggests that we describe harms to which a person is liable as narrowly proportionate, and harms to which a person is not liable, but that are nonetheless justified as the lesser evil, as widely proportionate. Similarly, we might think that that collateral harms to which a person has not consented are constrained by wide proportionality, just like harms to an innocent bystander. Such harms must be greatly outweighed by the good that one can achieve. It is, for example, permissible for Victim to render Bystander temporarily unconscious in order to save his life, as in *Attack*, but it is not permissible for him to kill Bystander to save his own life, as in *Collateral*. In contrast, collateral harms that a person consents to bear might be judged according to the narrow proportionality constraint. Just as one can forfeit rights by rendering oneself liable to harm, one can waive rights by consenting to be harmed. It might, therefore, be permissible to impose collateral harms that are narrowly proportionate, but not the lesser evil, on a person who has consented. (One difficult question is whether one may impose unlimited harm on a person who consents: for example, if Bystander in *Attack* agrees that Victim may kill her, would it be permissible for Victim to proceed?[17] However, we can set this aside here; our present aim is simply to show that whilst there are some collateral harms that may be inflicted only with consent, there are other harms, including non-trivial harms, that may be inflicted without consent.)

Of course, our ratio needs to be sensitive to the fact that there is often uncertainty about the success of a harmful action in securing the good. This uncertainty may be underpinning the very cautious approach to bystander risk exemplified by the Ordinary Life and Near Zero standards. But there is also uncertainty in the other cases we have considered: we don't know whether the police will catch the criminal, whether the paramedics will save a life, or whether defensive force will avert a threat. Thus, the uncertainty of securing the benefits of a trial does not undermine the appropriateness of employing a ratio account for trials, but merely shows that our account must take probabilities into consideration. The foregoing arguments suggest that, even granting the uncertainty regarding the success of trials, both the Ordinary Life and Near Zero standards are overly cautious, given the very great harms that successful medical trials prevent.

HARM TO BENEFICIARIES

The Beneficiary Principle and the Asymmetry Principle

The degrees of risk and harm to which a person may be exposed might also be affected by whether she stands to benefit from the risky action. For example, some philosophers argue that a rescuer may make the subject of a rescue bear the cost of her own rescue, provided that she is still an *ex ante* beneficiary. Consider *River*:

> *River*: Swimmer is drowning in a river. Passer-By can pull Swimmer out by either (a) breaking Passer-By's wrist, or (b) breaking Swimmer's wrist.

McMahan suggests that it is permissible for Passer-By to make Swimmer bear the cost of her own rescue in this case. Swimmer is significantly better off alive with a broken wrist than if she is left to drown. Call this claim – that one may make prospective victims bear the cost of their own rescue – the *Beneficiary Principle*.

The Beneficiary Principle looks especially plausible when the cost to the rescuer is so high that saving the victim is supererogatory. Say that pulling Swimmer out of the river requires either the loss of Passer-By's leg, or the loss of Swimmer's leg. Since Swimmer is still better off overall if Passer-By intervenes in a way that causes the loss of Swimmer's leg but prevents her death, it seems permissible for Passer-By to intervene by making Swimmer bear the loss of her leg. The Beneficiary Principle seems particularly pertinent to harms to bystanders in trials, since some bystanders will also be potential beneficiaries of the trial.

It would be impermissible, however, to impose the loss of a leg on a bystander for the sake of saving someone else from drowning. The bystander would simply be made worse off. This suggests another principle – which we can call the *Asymmetry Principle* – that holds that one may greater impose risks and harms on beneficiaries than one may impose on bystanders.

These principles imply that the permissibility of clinical trials might sometimes be localised – that is, it might be permissible to expose a local population to increased levels of risk as part of a trial from which they stand to benefit. This could include cases in which a disease affects some groups more than others, such as skin or breast cancer. It could also be relevant to trials on fairly universal diseases – such as renal cancer or Alzheimer's – if a particular local community is unlikely to be in a position to obtain the drugs if the trial is a success. For example, if the resultant treatment will be unaffordable to those in the local community, they may not benefit from the trial even if they are also affected by the disease.

Sometimes, of course, one might be an *ex ante*, but not an *ex post*, beneficiary of a risky action. Consider *Crocodile*:

> *Crocodile*: Swimmer is unconscious, and has a 90% chance of losing a hand to a nearby crocodile. Passer-By can fire a Taser into the water to stun the crocodile, but there's a 10% chance the shock will also paralyse Swimmer from the waist down.

Assume that partial paralysis is worse than the loss of a hand. It might nonetheless be permissible to use the Taser, given the different probabilities of each harm. Swimmer is still an *ex ante* beneficiary of the rescue attempt. Even if things go badly and she is worse off overall, she was still more likely to benefit. Likewise, it could be permissible to conduct trials that expose bystanders to risks of harms provided that they are *ex ante* beneficiaries, even if it turns out that they are in fact harmed rather than benefitted.

Harm and injustice

I think both the Asymmetry Principle and the Beneficiary Principle are *prima facie* plausible, and that being a beneficiary can be morally relevant to the distribution of risks. But both principles start to look less relevant when we move away from thinking about

natural dangers, like falling in rivers, and more about injustices that arise from wronging. Let me explain why, before outling the implications for clinical trials.

One implication of the Beneficiary Principle is that in a war of humanitarian intervention, intervening combatants may make the beneficiaries of the war – say, those people facing genocide – bear the costs of the interventions, shifting risk away from themselves onto the civilians. And, according to the Asymmetry Principle, the combatants may shift those risks much more onto members of the group facing genocide than onto other civilians.

But why would the fact that some people are the intended victims of a genocide weaken their rights against being collaterally killed compared to other people?[18] Would it not be more just to distribute the risks as evenly as possible? By way of comparison, imagine that a serial killer is attacking red-haired women, and extra policing is required to keep those women safe. It seems unlikely that we should pay for the policing by taxing red-haired women more. Rather, we should tax everyone more, since it is not the women's fault that they are at risk of being unjustly harmed.

On the face of it, it might seem as if this worry does not apply when we are trying to treat disease. Disease might seem more like a case of pulling someone out of a river – that is, a kind of natural phenomenon that is not really an injustice. Consider the malaria example I gave earlier. We might think that because people living in malaria-endemic areas stand to benefit so much more from malaria treatments, it therefore permissible to conduct trials that expose them to higher chances of infection. But of course, so much vulnerability to disease is an upshot of poverty, and poverty is by and large a result of injustice. This gives us reason to be skeptical about the moral significance of being a beneficiary when it comes to trials on poverty-related diseases. Exposing those who are vulnerable to disease as a result of injustice to greater degrees of risk is akin to imposing higher taxes on the red-haired women. But beneficiary status might matter for other sorts of disease (e.g. genetic disorders) that are not the result of injustice.

Refusal of consent

There is a further way in which being a beneficiary might be relevant to the permissibility of imposing risk and harming. In *Crocodile*, where Tasering the crocodile risks injuring Swimmer, I stipulated that Swimmer is unconscious. If, however, she was conscious and refused her consent, the rescue attempt would be impermissible. Exposing her to the more serious harm is impermissible if she would rather take the greater chance of the smaller harm. Similarly, should an affected bystander who might benefit from a trial refuse her consent to being exposed to risks of other harms, her refusal could make conducting the trial impermissible.

Beneficiary consent has even wider implications that are also of relevance for clinical trials. As Jonathan Parry has argued, it seems plausible that beneficiaries can refuse being rescued even if there is no concomitant risk that they will be harmed.[19] Call these people *pure beneficiaries*. For example, if Swimmer is a life-long campaigner for crocodile rights, she might refuse her consent to having the crocodile Tasered for her sake even if there's no

chance that the Taser will harm her. She might prefer to risk the loss of her hand than harm the crocodile.

Parry argues that we can explain the impermissibility of rescuing in such cases by recognising that one has control over one's interests, such that one gets to decide whether one's 'good' is used as a justification for harming. This affects what it is proportionate for rescuers to do. Consider *Pacifist Trolley*:

> *Pacifist Trolley*: A runaway trolley is heading to where it will kill five people. Onlooker can divert the trolley down a side-track, where it will kill Workman. Three of the five are lifelong pacifists, and shout that they do not want Workman killed for their sake.

Since three of the five refuse their consent to being rescued, Onlooker may weigh only the saving of the remaining two lives against the killing of the one. Killing one to save two is probably disproportionate and thus, Parry argues, Onlooker is required to let all five die.

If this is correct, it has important consequences for the role of bystanders in medical trials, even if those bystanders are not at risk of harm. Someone who is a pure beneficiary of someone else's being exposed to harm can refuse her consent to having the benefit to her included in the cost-benefit analysis of that trial. If enough potential beneficiaries refuse consent, the trial could be disproportionate. It matters a lot, then, whether we have identified beneficiaries – whether they are at risk of harm or not – or merely statistical beneficiaries. Statistical, unidentified beneficiaries are a bit like unconscious people – we permissibly assume consent when we think we will make them better off overall. But the same is not true of identified beneficiaries from whom we can obtain consent, and whose refusal of consent makes a difference to the cost-benefit analysis. Of course, we need to factor in benefits to future generations – it might be that these always outweigh the interests of the smaller number of living beneficiaries. But we cannot simply assume this – we often discount the interests of future people compared to existing people, and it might be that we should also do so in at least some medical trials.

CONCLUSION

This paper has explored some of the complex issues surrounding bystanders in medical trials. I have argued that the recently proposed threshold accounts of permissible risk imposition are implausible. Instead we should employ a ratio account, weighing the potential risks to bystanders against the potential benefits of the trials. Thinking about other permissible risk-imposing activities shows that it is sometimes permissible to impose significant risks of serious harms on some people for the sake of others. We should thus reject both the Ordinary Life and Near Zero standards for permissible risk imposition. Moreover, the risking of harm does not always require consent, even if the risk is high and/or the harm is serious. The reason we require consent from participants is that they risk being harmfully used as a means; bystanders risk being harmed as a side-effect, and non-consensual side-effect harms are easier to justify.

I also suggested that being an *ex ante* beneficiary of a trial can make it permissible to impose more risk on that person without her consent compared to someone who does not stand to benefit. However, I added that the significance of being a beneficiary can vary depending on the type of harm that we are trying to prevent – specifically, being a beneficiary of a risky action might not be morally significant if someone is vulnerable to harm as a result of injustice. Finally, I considered the role of refusing consent in establishing permissible levels of risk imposition. I suggested that pure beneficiaries are entitled to refuse consent for the use of their interests as a justification for a trial that imposes risks on other people. That, in particular, could affect the permissibility of conducting a wide range of clinical trials.

Acknowledgements:

I am grateful to participants at a workshop at the Harvard T. H. Chan School of Public Health for comments on an earlier draft. Special thanks to Nir Eyal and Lisa Holtzman for additional comments.

Funding: This contribution was enabled by NIAID grant R01AI114617 (HIV cure studies: risk, risk perception, and ethics), and Knut and Alice Wallenberg Academy Fellowship grant 1521101.

Biography

Helen Frowe is Professor of Practical Philosophy and Knut and Alice Wallenberg Academy Fellow at Stockholm University, where she directs the Stockholm Centre for the Ethics of War and Peace. Her books include *Defensive Killing* (OUP, 2014) and *The Ethics of War and Peace: An Introduction* (Routledge, 2011, 2015).

References

1. As Jonathan Kimmelman notes inKimmelman J (2005). Medical research, risk, and bystanders. IRB: Ethics and Human Research. 27(4), 1–6, p. 3.
2. The types of harms to which trials might expose bystanders vary. I focus here on risks of physical harm, such as the risk of life-threatening infections.
3. Notable exceptions includeKimmelman J (2007). Missing the forest: Further thoughts on the ethics of bystander risk in medical research. Cambridge Quarterly of Healthcare Ethics. 16, 483–490; [PubMed: 18018930] Hausman D (2007). Third-Party risks in research: Should IRBs address them?. IRB: Ethics and Human Research. 29(3), 1–5;Resnik D and Sharp R (2006). Protecting third parties in human subject research. IRB: Ethics and Human Research. 28(4), 1–7.
4. See also, Resnik & Sharp, op.cit. note 2, p. 4 – 5.
5. Battin MP, Francis LP, Jacobsen JA, & Smith CB (2009). The Ethics of Research in Infectious Disease: Experimenting on this Patient, Risking Harm to That One In Battin MP, Francis LP, Jacobsen JA, & Smith CB The Patient As Victim and Vector: Ethics and Infectious Disease (pp. 164–183).Oxford: Oxford University Press; Shah SK, Kimmelman J, Lyerly AD, Lynch HF, McCutchan F, Miller FG, … Zorilla C (2017). Ethical Considerations For Zika Virus Human Challenge Trials: Report and Recommendations. National Insitute of Allergy and Infectious Diseases.
6. I remain neutral here on what counts as a relevant benefit.
7. Battin et al., op. cit. note 5, p. 5
8. Shah et al., op. cit. note 5, p. 19
9. For example, the US federal regulations stipulate that “risks to subjects are reasonable in relation to anticipated benefits, if any, to subjects, and the importance of the knowledge that may reasonably be expected to result…” (CFR446.111 Criteria for IRB approval of research).Eyal N (2018) See also, e.g.What can the lived experience of participating in risky HIV cure-related studies establish?. Journal of Medical Ethics. 44, 277–278; [PubMed: 29321219] Emanuel EJ, Wendler D, & Grady C

(2000) What makes clinical research ethical?. Journal of the American Medical Association. 283, 2701–11. [PubMed: 10819955]

10. The original trolley case was proposed by Foot P (1967). 'The Problem of Abortion and the Doctrine of Double Effect', Oxford Review. 5, 1–7.The variation of the bystander at the switch comes fromThomson JJ (1985) The Trolley Problem. Yale Law Journal, 94(6),1395–1415), p. 1397.
11. See, e.g., Hallborg R Jr. (1997). Comparing Harms: The Lesser-Evil Defense and the Trolley Problem. Legal Theory. 3(4), 291–316;Kamm FM (2007) Intricate Ethics: Rights, Responsibilities and Permissible Harm. Oxford: Oxford University Press;McMahan J (2014) Self-Defence Against Justified Threateners In Frowe H & Lang G How We Fight: Ethics in War Oxford: OUP, pp. 104–137, p. 105
12. See, e.g. Wendler D, & Emanuel EJ, (2005). What is a 'minor' increase over minimal risk?. Journal of Pediatrics. 147(5), 575–578.
13. It may also be sensitive to, e.g. whether someone is culpable with respect to some threat. But since we're interested only in innocent people here, those complexities need not concern us.
14. See, e.g. Frowe H (2014). Defensive Killing. Oxford: Oxford University Press;Tadros V (2016) Dimensions of Intentions: Ways of Killing in War In Lazar S& Frowe H Oxford Handbook of Ethics of War (401–417). New York: Oxford University Press;Tadros V (2011). The Ends of Harm: The Moral Foundations of Criminal Law. Oxford: Oxofrd University Press, pp. 117–122.
15. See, e.g. Rodin D (2011). Justifying Harm. Ethics. 112(1), 74–110, p. 97;Frowe H (2018). Why We're Required to Turn the Trolley: On Lesser-Evil Justifications for Harming. The Philosophical Quarterly. 68(272), 460–480.
16. Since I'm drawing on McMahan's wider view, I'll stick with the deontological reading of the case. But consequentialists could also believe that, e.g. it is better that culpable people due than innocent people, of course.
17. One possibility is that consent allows us to simply compare the outcomes impartially: since the loss of a leg is impartially better than death, I may cause the loss of a leg of a consenting bystander to save my own life, even though it would be impermissible to do this without her consent. But it would plausibly be impermissible to kill a bystander for the sake of preventing the loss of my leg, even if she consented, since the outcome would be impartially worse.
18. See, e.g. Christie L (2018). Distributing Death in Humanitarian Interventions In Jenkins R, Robillard M, & Strawser BJ Who Should Die? The Ethics of Killing in War (186–201). New York: Oxford University Press.
19. Parry J (2017). Defensive Harm, Consent and Intervention. Philosophy and Public Affairs. 45(4), 356–396