Is There an Ethical Upper Limit on Risks to Study Participants?

Nir Eyal*, Center for Population-Level Bioethics, Department of Philosophy, Rutgers University and Department of Health Behavior, Society and Policy, Rutgers School of Public Health

*Corresponding author: Nir Eyal, Center for Population-Level Bioethics, Rutgers University, 112 Paterson St, Rm 400, New Brunswick, NJ 08901, USA. Tel.: +1 609 933 2840; Fax: +1 732 932 1253; Email: nir.eyal@rutgers.edu

Are some risks to study participants too much, no matter how valuable the study is for society? This article answers in the negative.

Caps on Risk to Study Participants

Clinical research ethics requires, among other things, a favorable balance between the risks to a study's participants and the expected social value of that study, as well as keeping such risks to those strictly necessary for realizing that social value through a valid study (Emanuel *et al.*, 2000; HHS, 2009; Rid and Wendler, 2011; Wikler, 2017). Yet some authors and guidelines warn that certain risks are too high, such that even studies with the highest promise for social value, say, urgent research like the one described above, could not exonerate them. In their view, there is an upper limit or (as I shall put it) cap on study risks to participants.

Some say, e.g. that when a high probability of death or severe injury to a study participant from an experiment is predicted, the experiment cannot be ethical, no matter what; any experiment involving such risk remains illegitimate, even if society stands to gain tremendously from conducting it, the quality of consent is high, risks are being minimized within bounds of study validity, review is independent and proper, and so forth. What would rule out such a study is an alleged cap on how much individual participants' health and welfare can ever be compromised on the path to scientific, medical and humanitarian progress.

Or so the view goes. This article argues that this view is false. My thesis is that, when the research is important enough, and everything else is done right, there is no upper limit on the risks that studies may legitimately visit upon their participants.

Let me start by presenting the notion of caps on permissible risk to study participants. I shall then make a few points against the existence of such caps, and then,

answer attempts to defend them. An Appendix will propose a number of distinctions between caps on risk to participants.

Certain challenge trials in public health emergency circumstances offer dilemmas about transgressing limits of risk to participants. The issue arose recently because it is urgent to roll out proven vaccines against the novel coronavirus. A human challenge trial was proposed for testing vaccine efficacy more rapidly (Eyal *et al.*, 2020; Plotkin and Caplan, 2020; WHO Working Group for Guidance on Human Challenge Studies in COVID-19, 2020). In challenge studies, participants are randomized to receive either the vaccine being tested or placebo, and are deliberately exposed to the pathogen. Soon thereafter, it becomes plain whether the vaccine beats the placebo in preventing infection and other outcomes of interest, or not. Opponents countered that exposing volunteers to a deadly virus for which no cure exists is too dangerous (Dawson *et al.*, 2020; Shah *et al.*, 2020). I shall take it for granted that regular coronavirus challenge trials, whose mortal peril for healthy volunteers in their twenties (per current proposals) has been estimated (e.g. by Jamrozik and Selgelid, 2020) at levels that are lower than those of live kidney donation, are acceptably risky. But it is true that there are ways to either accelerate such challenge trials further or make their findings more easily generalizable to a broad population of target vaccine users, which would be much riskier. One is to accelerate the process of viral dose confirmation necessary before the challenge trial can start, at some expense to the safety of dose escalation volunteers.¹ Another is to recruit not only young and healthy volunteers but also older or sicker volunteers, so as to increase the generalizability of the findings to patients with risk factors for severe

COVID upon infection (Shah *et al.*, 2020). Such tweaks to the safety-conscious designs thus far proposed would serve global public health, by expediting rollout for a vaccine that is confirmed safe and efficacious all users, but they would be riskier for participants, and potentially transgress alleged norms on maximal risk to participants of medical research.

Or take a more hypothetical yet unfortunately realistic case, of a disease with far greater virulence than COVID-19. A strain of coronavirus as infectious as seasonal flu yet far more lethal than COVID—say, a form of Middle East Respiratory Syndrome (MERS)—breaks out in a civil war zone. Either standard containment strategies or investigating the efficacy of a safety-tested vaccine candidate in the field are impossible in that dangerous zone, which people flee to many international destinations. This arguably realistic scenario creates a dilemma. Either defer vaccine efficacy testing until the outbreak reaches other areas, risking millions of lives; or invite a small number of volunteers now to receive the candidate vaccine and then be exposed to the relevant MERS strain vs. placebo control. The latter study would be extremely risky for volunteers in either arm (because the strain is so lethal and the candidate vaccine, merely experimental and offered only in one). But the study could also accelerate effective response and be extremely socially valuable.

A decade ago, Miller and Joffe (2009: 445) raised an excellent question: ‘Is there a maximum level of net risks to consenting research subjects that can be justified by the potential social benefits from a particular scientific investigation?’. As they point out, the Nuremberg Code, promulgated in the wake of Nazi wartime experiments, which were abusive on many fronts, stipulates a limit on permissible research risks. The Code decrees, ‘No experiment should be conducted where there is an *a priori* reason to believe that death or disabling injury will occur’ (International Military Tribunal, 1947: 182).

To characterize what levels of risk such caps on study risk forbid, Miller and Joffe proffer as a comparator the level of risk to donors of live organs that is typically considered acceptable. They do not commit, however, to the existence of upper limits or, as I shall put it, to ‘caps’ on risk to participants in clinical studies. Nor are they entirely clear on how much the caps coincide with standards for permitting organ donation.

Other leading bioethicists endorse caps on research risk to participants more definitely. For example, Alex London urges researchers and oversight bodies ‘to ensure that the incremental risks to the basic interests of research participants that are associated with purely research-related elements of an investigation are not

greater than the incremental risks to the basic interests of others in the community who work on a routine basis to advance the common good’ such as, e.g. fire fighters (London, 2009: 1193). And elsewhere in the same article, London endorses a standard for the upper limit to permissible risk in those studies whose consenting participants are healthy. That standard allows even less incremental risk in any trial arm (compared to the other arm); at that point, the cap seems to be at *zero* incremental net (comparative) risk from the medical community’s standpoint (London, 2009: 1191). Whichever standard added net risk from study participation somewhere. Notably, in London’s view, these caps govern not only ordinary research but also studies of great urgency and social value, e.g. certain humanitarian studies; for London, caps on risk only become more important during public health crises, where he traces a general tendency to overburden vulnerable study participants—a tendency that, he cautions, must be resisted (London, 2009: 1201–1202).

Annette Rid and David Wendler are also clearly committed to the existence of caps on risk in clinical studies. In their words, ‘there are levels of net risks to individual participants that cannot be justified by even tremendous social value. If the cumulative net risks in the study clearly exceed the general limits of acceptable research risk, the study should be rejected’ (Rid and Wendler, 2011: 162).

David Resnik agrees, and proposes another concrete cap on risk of clinical studies for healthy volunteers, namely, ‘a 1 per cent standard: healthy volunteers in biomedical research should not be exposed to a greater than 1 per cent chance of serious harm, such as death, permanent disability or severe injury or illness’ (Resnik and Sharp, 2006: 147). While Resnik claims that ‘this standard... is not an absolute rule’, he adds that even when ‘compelling public health or social problems’ justify going above that number, applications must show ‘that the risk of serious harm is only slightly more than 1 per cent’ (Resnik and Sharp, 2006: 147). So it seems fair to conclude that, while Resnik does not state what the absolute cap is in studies of exceptionally high social value on healthy volunteers, he holds that such a cap exists even in such studies. For example, Resnik would probably reject a study that carried a 3 per cent risk of seriously harming healthy study participants, no matter how socially valuable the study may be, because 3 per cent is more than ‘slightly more than 1 per cent’.

A thorough review by Rid reveals that when no benefit for participants is expected, many research ethics guidelines and regulations espouse caps on the risk to study participants.² And when waivers or modifications of

informed consent are necessary for methodological reasons (e.g. research involving deception) or when the research could not practicably be carried out without a waiver (e.g. medical records research), many regulations and guidelines require that the risks to participants be no greater than minimal—which can be read as very low caps on the risk to them (Rid, 2014: 64).

Indeed, the list of authors and documents endorsing caps on acceptable risk to participants may be even more extensive. Recall that Miller and Joffe, Rid and Wendler, and others write approvingly of caps on *net* risk to study participants. Obviously, the net risk to a participant can remain very high even when some benefit to her is present (Rid, 2014). Surely the case for caps on risk to participants of ‘non-beneficial’ research can be used to motivate such caps in somewhat-beneficial research. For example, an experimental substance may be likely to relieve a participant of her longtime chronic pain but also to trigger her death a few months later; studying that substance on people may count as beneficial or therapeutic research inasmuch as overcoming pain is a therapeutic benefit; but the *net* risk involved is so high that many would consider this study so harmful on balance as to be illegitimate. And they would forbid it even if the study had tremendous social value, say, for proof of concept of new interventions to eliminate severe pain in vast numbers of dying patients. They would be espousing a cap on risk in a study that may benefit unhealthy participants.

In short, many authors and guidelines would endorse, or are committed to endorsing, caps on acceptable risk to study participants. While their position takes on different forms and degrees (see Appendix), the common denominator can be formulated as the position:

Caps on acceptable risk: There is a level or a kind of risk such that visiting it upon study participants for study purposes alone is never permitted, no matter how socially important the study is, and how much the study is otherwise legitimate, e.g. study risks have been minimized, consent is highly informed and so forth.

The following two figures illustrate the difference between a simple balance between risks to participants and social value on the one hand, and a *caps on acceptable risk* approach on the other. In the two models depicted in Figure 1, there is no such cap. Added risks to participants and the social value that can warrant them stand in either a simple ratio (in the model on the right) or a more sophisticated relation (in the one on the left). Both the model on the right and the one on the left of Figure 1 accept some studies’ risk-benefit profiles (examples are

studies a and d) and reject others’ (b). The difference between the models to the right and to the left of the figure comes into relief in other studies, which one model accepts and the other rejects (c). Both models’ failure to endorse caps on net risk to participants is shown by both models’ approval of a study carrying extraordinary risk to study participants, but also tremendous social value (d).

In Figure 2, there is a cap on the risk to participants (represented by a dotted line). Studies whose risk for participants exceeds that cap (d) are rejected. The area to the right of the dotted line in either model of that figure continues upwards with no limits. Put differently, no finite level of social value would legitimize any risk exceeding that cap. When balancing overall reason to limit risks to participants vs. the overall reason to generate social value, beyond that cap the reason to limit risks to participants is weighted by infinity. The model on the left shows that this need not involve an abrupt break.

Intentionally, these figures do not specify what is meant by ‘risk to participants’ and what makes that risk worse, e.g. whether the risk is net or not; whether it represents the combined risk to all participants, or the risk to highest-risk participants, or to the average participant; and whether the probability of harm weighs as much as the size of that harm. Importantly, the below argument against caps to acceptable risk for participants can stand, on any of these specifications.

A further clarification: There is another way to conceive of caps on risks to participants, which is not well-represented by such figures. Suppose that the social value of studies, or the strength of the reason to conduct them given their value, simply could not exceed a certain level. That would also cap how much risk to participants it is worth accepting in return for attaining that value, but without the assumption, criticized herein, that even infinitely valuable studies could not excuse high-enough risks to participants. Rather, on that alternative approach, the reason for caps is that no study ever has infinite social value. To put the point more generally and philosophically, when one value or practical reason is ‘superior’ to another, such that the first one overrides the other no matter how stringent that other value or reason is, that need not show that infinite weights are properly assigned to the one; it can instead show that the overridden value or reason is barred from ever passing a certain modest ceiling, which the overriding one is able to cross.³ For medical experiments, however, this alternative possibility can be largely assumed away.⁴ There is usually no modest ceiling on the social value of medical studies.⁵ Clearly, some experiments are tremendously

Figure 1. Two ways to balance risks to participants and social value, without caps on the risk to participants.

Figure 2. Two ways to balance risks to participants and social value, with caps on the risk to participants.

valuable, and vastly more valuable than ones of ordinary value. That difference in social value is reflected, e.g. in the appropriately different limits on public funding for these respective types of experiment. One reason why a vaccine efficacy study with high promise to halt an otherwise catastrophic pandemic outbreak should be budgeted very generously (as the USA is currently budgeting coronavirus vaccine studies) is the extreme social value of such a study, if properly done and used. The worse the outbreak that the study is likely to stop, the greater the imprudence of trying to save money by not conducting

it. In philosophical jargon, we might summarize this as follows. It is true the superiority of one reason or value to another does not always assume that in their balance, the former has an infinite weight. But in our context, this complication does not arise. An infinite weight is the only plausible interpretation of a fundamental ethical cap on risk in medical studies and the implied superiority of protecting participant from great risks over augmenting social value.

A final clarification: *Caps on acceptable risk* covers public health emergency circumstances as well. As we

saw, authors like London are emphatic that caps apply in such emergency circumstances, and the regulations cited do not make exceptions for emergencies. At least as I have phrased *Caps on acceptable risk*, it is not a mere recommended guidance or recommended law, which might have been conjoined with the clarification or secret plan that different instructions would take effect in emergency. Instead, *Caps on acceptable risk* is phrased as a right-making ethical standard with the purport to cover all circumstances.

As we saw, many authors, guidelines and regulations in research ethics are committed to or have explicitly endorsed some variant of *caps on acceptable risk*. Indeed, it is harder to find authors or documents who accept a risk-benefit requirement in research while explicitly rejecting caps on risk to participants.⁶ Nonetheless, this article will argue against the possibility of caps on risk to study participants.

Why WE Should Resist Caps on Risk to Study Participants

There are four independent considerations against caps on risk to study participants in research.

The Rejection of Moral Absolutism

The overwhelming majority of contemporary philosophical ethicists rejects absolutism about unconditional moral rules. That overwhelming majority includes all consequentialists and most anti-consequentialists. So-called threshold non-consequentialists contend that small gains in overall value do not justify infringement of moral thresholds, but agree that extreme gains justify their infringement (Nagel, 1972; Nozick, 1974; Moore, 1989; Thomson, 1990; Kamm, 2001).⁷ In fact, the standard example they cite is the extreme social value of overcoming public health emergencies, presumably including high-risk studies to thwart epidemic and pandemic threats. While many contemporary Kantians reject such thresholds and are in that sense absolutist, they are absolutist only about conditionally phrased act types, which would accommodate different procedures for public health emergencies. These Kantians would presumably reject unconditional caps on risk to participants. Unconditional absolutist major ethicists can be counted on the fingers of one hand.

Thus, both consequentialists and the vast majority of contemporary anti-consequentialists would need to reject caps on risk to participants, given these caps' unconditional and absolutist nature. Certainly to reject these

caps does not commit one to crude utilitarianism on research ethics. In Figure 1 above, the complex approach on the left is anti-utilitarian and anti-caps. Even the approach on the right in that figure can only be utilitarian if the ratio between risks to participants and the cumulative social value that is thought to legitimize them equals 1 per person affected; choose any other ratio and you give extra weight either to study participants or to the people who need novel interventions to fight their current and future ailments—against utilitarian advice. A position that rejects *caps on acceptable risks* can also remain anti-utilitarian if it opposes coercive studies, if it prioritizes the disadvantaged, and more.

If *caps on acceptable risks* purported to cover only non-emergency circumstances, it might be thought to avoid absolutism. It would instead pertain only to modest-value studies in which caps would not rule out a study with tremendous social value. However, we have seen that *caps on acceptable risks* purports to apply to all studies.⁸

Participants' Autonomous Authorization

Second, participation in high-risk studies can remain consensual. By contrast, philosophical discussions of whether it is OK to sacrifice one innocent for the sake of many (*Trolley, Jim and the Indians, Transplant, Georgia Jury* and the like) are nearly always about an innocent who is loath to be sacrificed, or is cheated. The same goes to nearly all historical medical studies that abusively sacrificed individuals for others' sakes. The consent of capacitated adults to being treated in this or in that way both legitimates treating them in that way and may make a third party's blocking such treatment into overreach. There is a world of difference between interfering to stop adults from interacting in a way forced by one of them on the other, or done under manipulation, and interfering to stop adults from interacting in perfectly consensual ways, whose externalities for others are net positive. The former is usually a duty. The latter is usually forbidden.

An opponent may retort that consent to a high medical risk cannot be valid. But high risk for study participants in medical and direct terms does not mean that their participation is necessarily unfree, uninformed, incompetent or irrational. For one thing, they may prospectively benefit from participation, in indirect medical and in non-medical terms (Eyal, 2017, 2020). More importantly, in the service of collective good, it can also be perfectly rational to knowingly jeopardize one's overall health and welfare a lot (Agrawal and Emanuel, 2006);). Even if seriously risking oneself for a trivial cause would

have been irrational (contrast with [Buchak, 2017](#)), doing so to save many people from substantial personal risk, or, in some scenarios, greater personal risk than one is taking on, can surely be rational and laudable (Parfit, 2011: 130–140).

An opponent who denies the possibility of free and informed capacitated consent to take a large medical risk may instead be making a psychological claim. She may be arguing that, as mere psychological necessity, fully free, informed and capacitated human beings would never join very risky studies, so consent to any such studies must be flawed (compare [Steel, 2019](#): 213). But people vary, and psychological necessities are few or nil. Considering humanity's large population, surely a few dozen volunteers with full understanding and competence would accept grave risks to help save scores. Challenge studies require only a few dozen volunteers, and in global public health emergencies, would be prudent to transport them to study sites from even far afield ([Eyal, 2020](#)). Surely somewhere, one could find a few dozens of special people.

The correct response to the risk of unfree or uninformed consent is not to prevent people from voluntarily taking high risks. It is to take measures to ensure full freedom and full comprehension of the high risks. I say that as someone who sometimes supports either soft or hard paternalism; in this case, however, overall good is not served by stopping individuals from joining risky studies—because the relevant studies do great good.

Real Life

As a final strike against caps on risk to study participants, these caps have usually been proffered without picturing seriously the extreme scenarios that lend initial support to high-risk studies. Picture such scenarios, and these caps start losing appeal. This is what gradually happens to many of us as a public health disaster unfolds, and we are considering the permissibility of a risky study that could help end it. Outside disaster, the gut feeling or intuition is initially strongly negative, but further imagination of the full societal implications of unbridled disaster tends to open us more to accepting risky studies. Even in the imagined MERS case above, once the full global implications of the decision are vividly imagined, I believe that it is no longer as intuitive that willing individuals must be barred from volunteering to take on the trial's severe personal risks.

In the same vein, consider spring 2020 warnings by bioethicists interviewed in the popular and scientific press about expediting coronavirus vaccine development through challenge studies. Almost none were

reported to attempt quantifying the cumulative harm from failure to expedite vaccine development and roll-out and only then explaining why the risks to participants should nevertheless weigh more than the cumulative mortality and morbidity from, e.g. COVID, neglected medical services, halted economic development and other sequelae of the delay they effectively proposed. One interpretation is that most of these bioethicists either never seriously imagined the enormous human toll from their recommendations, or wanted to suppress such imagination among readers whom they were trying to convince against challenge studies. As a result, their analyses expounded one side of the equation only.

As a final illustration, Walter Reed's Yellow Fever challenge study was high risk and killed some study volunteers but identified the mosquito vector for Yellow Fever, thus preventing an untold number of deaths from 'the scourge of the South' and from other mosquito-borne killer diseases. Few would call it a flagrant violation of moral edicts of research, however. Again, when the full humanitarian costs of caps on risky research are given our full attention, caps on risk to participants lose much of their intuitive appeal.

Admittedly, in considering such real-life scenarios, our intuitions may be clouded by self-interest, or we may just be too scared to think well. But the opposite is also true. Outside disasters, disaster scenarios often seem to people as though they could never materialize—although they could. Indeed, even as disasters unfold, optimism bias and problems imagining exponentially growing fatalities can interfere with grasping the full value of expensive preventative strategies. It is possible that the same gleeful optimism that regularly thwarts true societal efforts to prevent and prepare for most disasters also thwarts our willingness to accept hard ethical decisions about them. It makes it more convenient to stick to lofty ideas about a limit on the permissible risk to individuals in tertiary and in crucial research alike.

To summarize my three arguments, even for the most stringent rules against risking and harming innocent study participants, there will always be a study that has such extremely high social value as to warrant infringement of these rules; participants can and should provide free and informed consent to the related risk, thereby making any applicable moral rule against jeopardizing study participants significantly less stringent. The relaxation of the rule makes it easier for the study's social value to warrant the risks to participants on balance. Our intuitions suggest as much once we imagine being immersed in such a real-life situation, and such

immersion is a vantage point for the evaluation of moral intuitions.

Responses to Arguments for Caps on Risk to Study Participants

Prohibitions on high-risk studies are typically stipulated and rarely defended with a fully fledged argument. As we shall see, the few existing or conceivable attempts to shore them up fall flat.

The Uncertainty of Studies' Social Value

Joffe and Miller (*op. cit.* p. 448, note 15) try to motivate a cap on risk to participants through the thought that, in real-life examples, the social value of studies is never certain. Indeed, that uncertainty is typically the point of conducting such studies. And even success in all study phases would not secure social value. The latter can require rapid approval, mass-scale production and judicious rollout of the intervention being tested as well, and those need not materialize.

Nonetheless, I would insist, when social stakes are high enough—say, in worst-case scenarios like the MERS scenario above, the prospective social value of the study may remain extremely high, uncertainty notwithstanding. Especially when there is some basic assurance about proficiency and coordination between the relevant investigators, approval agency, manufacturer and health agencies, and concrete plans for later rollout, expected social value could remain high enough to make the extreme risks worthwhile.

Here is a back of the envelope calculation. The Spanish flu of 1918 infected an estimated 500 million people worldwide. Part of the reason for the enormous number of infections was the lack of technological advances like vaccines (US Centers for Disease Control, 2018). The pandemic MERS case I proposed was meant as a realistic worst-case scenario before relevant vaccines are developed. So it is conservative to suppose that in our more populous and more connected world, absent a vaccine (i.e. relying merely on therapies, good personal hygiene, isolation, quarantine and closures), there would be a 10th of the number of infections, namely, 50 million. If on the other hand a safe and efficacious vaccine is approved, produced and rolled out on time, some of these infections would be prevented. A conservative estimate is that it would prevent only a 10th, namely, 5 million infections. Because so far, 'approximately 35 per cent of reported patients with MERS have died' (WHO, 2018), to prevent those infections would prevent

approximately 1.75 million deaths (assuming non-trivially a similar fatality rate).

How much would the challenge study being considered facilitate preventing these deaths? Challenge studies rarely suffice for vaccine approval (Shah *et al.*, 2017) but they can substantially accelerate approval (Marston *et al.*, 2016; Shah *et al.*, 2017; Vannice *et al.*, 2019; Plotkin and Caplan, 2020; WHO Working Group for Guidance on Human Challenge Studies in COVID-19, 2020). While conventional efficacy studies more regularly lead to approval, only a half or two-thirds of vaccine candidates that undergo any efficacy testing reach approval (Pronker *et al.*, 2013; Hay *et al.*, 2014). After all, efficacy trials can reveal insufficient efficacy. And notwithstanding the special efforts that the dire situation may spur, a vaccine found to be safe and efficacious need not be approved, produced and rolled out on time to make this big difference. Taking all that on board, let us conservatively assume that the acceleration spurred by this challenge study would incrementally boost global population health only one-twentieth as much as finding a vaccine candidate safe and efficacious and then immediately approving and rolling it out widely and expeditiously. Since we estimated the latter as equivalent to preventing 1.75 million deaths, the marginal contribution to global population health of this MERS vaccine efficacy challenge study in the circumstances would be that of preventing $1.75 \text{ million} / 20 = 87,500$ deaths.

Does that justify the acute risk to study volunteers? Challenge studies 'only involve a few dozen volunteers' (Cohen, 2016). Assuming 150 volunteers, half of whom are placebo controls, and a pessimistic scenario in which this safety-tested vaccine turns out to be wholly inefficacious, such a study would take around 52 volunteer lives (35 per cent). But somewhat more optimistic scenarios, in which the vaccine candidate is found to be efficacious, and the same study takes fewer volunteer lives (26 volunteers in the placebo arm and 10, say, in the active arm, so 36 overall), are equally likely (or more likely: see Pronker *et al.*, 2013; Hay *et al.*, 2014: n. 40). We shall put the respective probabilities of the optimistic and the pessimistic scenarios at 0.5. While risking so many volunteer deaths from active investigational intervention remains absolutely daunting, the extreme circumstances arguably call for an extreme response, and $87,500 / (52 \times 0.5 + 36 \times 0.5) = 1988$ is an incredibly high ratio. The challenge trial, in other words, prevents 1988 innocent deaths for every innocent it kills (and later we shall note that the moral significance of the consent of the latter to be placed at risk).

By comparison, in Williams's (1973: 98–100) famous example of *Jim and the Indians*, where Williams seems to hold that it is not 'obvious' (pp. 99, 117) whether to kill one to prevent many killings, the ratio is 20 lives saved by one active killing. One introduction to ethics repeatedly illustrates deontological 'threshold' views as clearly permitting infringements when they would save 100 lives through one active killing (Kagan, 1998: 79, 82). While these philosophers' examples were only meant as illustrations, they may give us a sense of what specific numbers non-absolutist philosophers have in mind. Presumably, therefore, most non-consequentialists who stop deeming infringements of a right impermissible on balance beyond a certain 'threshold' of value gain would accept this challenge study. After all, 1988 lives are nearly 1000 times Williams's 20 lives figure and nearly 200 times the textbook's 100 lives threshold. The vast social value of the MERS challenge study makes the relevant killing permissible on balance, and certainly not so evidently impermissible as to require special regulatory prohibitions.⁹

To be fair, some philosophers' examples involve additional *prima facie* reasons to act transgressively, beyond the expected social value of that transgression, including *prima facie* reasons that are absent from our MERS example. For example, in *Jim and the Indians*, the 20 potential victims are well identified and their deaths would be active, wrongful and relatively certain, much more than typical victims of future MERS infections. So Kyungdo Lee suggested¹⁰ that these factors may sufficiently explain Williams's dilemma without thereby licensing risk to study participants in the MERS example. One might add that in some of the philosophers' examples, there are also non-consequentialist factors undermining the case against the transgression (in *Jim and the Indians*, perhaps the fact that the person whom one would kill is going to die anyhow because he is among the 20). Yet some of the philosophers' examples generate similar intuitive responses even in variants free from these complications (e.g. in *Jim and the Indians*, I think the immediate intuitions are fairly stable if the 20 who stand to be killed if one avoids killing one are unidentified, or if the one is not among the 20). Furthermore, some philosophers' examples also involve *prima facie* reasons to *avoid* the transgression that do not obtain in our MERS example, either (e.g. in *Jim and the Indians*, absolute certainty of killing, complicity in an evil plan)—so the threshold against the transgression may actually be *lower* than I have so far assumed. And even if in the philosophers' examples all the *prima facie* reasons to avoid the transgression are sound, they would presumably forbid a value gain that is a few times larger than

in the philosophers' examples; yet our MERS case involves a value gain that is 100–200 times greater than the ones discussed in the philosophers' examples.

A Principle of Equality

London (2009: 1191) grounds caps on risks to participants in a principle of equality, according to which,

as a necessary condition for ethical permissibility, research with human subjects must be designed and carried out so as not to undermine the standing of research participants as the moral and political equals of their compatriots, by either knowingly compromising their basic interests or showing unequal concern for their basic interests and the interests of the people the research is intended to serve.

London is surely right that in a sense, research that is especially risky for study participants is 'knowingly compromising their basic interests'. But in that sense, so does any research with net-negative prospects for participants, or at least some chance of causing lasting serious disability or death, including most toxicity studies on healthy volunteers, where for each participant the chance of a therapeutic effect is much lower than the chance of medical harm, and there is a chance of serious medical harm that surely would compromise basic interests. Yet such toxicity studies are widely accepted and considered consonant with human equality. While one could argue that people have a weighty basic interest in autonomy that legitimates letting them enter toxicity studies, the same could be said about high risk–high social value studies. There are also plenty of examples from outside human subjects research in which societies appropriately and legitimately knowingly risk or harm the basic interests of some, with or without their consent. We ration scarce basic resources, redistribute meager incomes to share them more fairly at times of need, and much else. In addition, London (2009: 1187) himself acknowledges that in a public health emergency, 'it seems reasonable that [community members] should be permitted to bear some affirmative risks to themselves in order to help their compatriots'. If some affirmative risks are compatible with maintaining their basic interests, and with his principle of equality, I fail to see why large affirmative risks to these same community members are incompatible with them—again, so long as the communal need is large enough.

Public Trust

London gives a further argument for caps on risks to participants. He is concerned that very risky research

may undermine public ‘trust’ in public health or in researchers at a time of crisis when public cooperation is needed. Risky research jeopardizes ‘the willingness of community members to believe the information that they receive from basic social and governmental institutions, to rely on and comply with their instructions, and to provide various forms of cooperation and support for their efforts’ (London, 2009: 1174; Dawson *et al.*, 2020).

All that may be true, but highly valuable experiments can also *promote* trust, by showing the competence of medical and public health institutions and the power of, e.g. vaccines. Such experiments may also make so much headway against a horrendous disease as to warrant on balance risk of some decline in public trust. For example, if investigators in high-risk studies were always participants as well, public trust would probably increase somewhat (Jonas, 1969), but requiring that might stifle research enough that it is not worth a moderate increase in trust.

The question what would promote public trust (or the most relevant subcategory in the context—trust that vaccines work; trust that researchers are scientifically competent; trust that researchers behave ethically, and so forth) is highly complex. It is also an empirical question in public health risk communication, a subdiscipline that most normative thinkers lack training in. To maintain public trust in researchers’ ethical conduct, perhaps risky studies should be done with manifest rigorous review; openness to direct public scrutiny; and media interviews with dedicated and well-informed participants, and with independent health experts and trusted religious and community leaders who push the (correct, if the study is ethical) message that less risky paths would have probably foundered. But we should resist a recent trend among bioethicists to anchor ethical recommendations in what would allegedly preserve public trust—amateur speculation over a matter of great psychological and sociological complexity.

Elsewhere London expresses openness to the possibility that his approach will fail on occasion to save trust, but writes as though the goal is not to maintain public trust but to ‘*merit* such trust’ (London, *op. cit.* note 17, pp. 1199–2000—original italics, shifting back from predictive foundations to normative ones). However, when a very risky study is otherwise perfectly permissible, conducting it would arguably maintain the authorities’ worthiness of trust. *Ex hypothesis*, they will have done nothing that is otherwise wrong! London may also believe that in pluralistic societies, controversial measures like high-risk studies are simply bound to diminish institutional trust too much, because he writes:

it seems reasonable to think that members of a pluralistic community will disagree about the nature and extent of the risks that social institutions can offer to community members, even in the service of noble ends. . . . the strains on institutional trust . . . will likely only be exacerbated if disagreement about the latter issue erupts in the context of a public health emergency (London, 2009: 1187).

But the same could be said about many correct responses to public health emergencies. Broad acceptability is a desideratum, yet there are better ways to maintain trust than to avoid responses to emergencies that would upset some, wholesale. Consider unpopular social distancing, safe burials, slightly invasive surveillance and contact tracing, careful quarantine, rollout of safe and efficacious vaccines that some irrationally oppose, and refusing ‘compassionate use’ that would invalidate a promising vaccine study. These measures stand some chance of aggravating some stakeholders but that does not mean that we should simply avoid them. We ought to at least try to educate the public about their importance. The same goes for emergency research that posits high risks to participants and may aggravate a subset of the population—in the absence of a strong independent argument against it.

The Primacy of Individuals

One natural thought is that caps on risk to individuals capture the prevalent ethos in research ethics that individuals come before collectives. Even the greatest social value cannot permit full compromise of an individual.

Whether or not individuals have such primacy, a research ethicist rejecting caps on risk to individuals, can espouse and express primacy to individuals, or to individual participants, in alternative ways. For example, she may hold that only very large social benefits justify high risks to individual participants. If she is a contractualist, she may hold that only large social benefits that include rescuing at least one other individual from commensurate or higher risks or from commensurate or greater harms can justify high risks to individual participants. If she is a libertarian, she may decree that the consent of the individual to any (large) risks for her is key. In short, there are many ways to express the alleged primacy of individuals without committing to caps on risk to participants.

History

Some of the support for caps on study risk probably stems from the prevalence of very risky studies among

the most unethical studies in the history of medical research. We should therefore make a point of simple logic. The fact that many (or all) highly unethical studies contained a certain element, e.g. very high risk to participants, does not show that every study containing that element is unethical. Likewise, all unethical studies were studies, yet not all studies were unethical.

Potential for Abuse

A final argument for a cap on risks to participants is concern about potential abuse. Regulatory permission to expose participants to high risks in extremely valuable studies increases the likelihood that participants would, abusively, be subjected to very high risk in modest-value studies as well. It may also increase the risk of exposure to high risks without fully informed consent.¹¹ It may at least increase the likelihood that potential funders and participants, worried that this might happen, would refrain from funding or from joining studies. Since extremely valuable studies are rare, this increased likelihood of abuse and of stifled research matters more than being able to respond optimally in extreme cases.

Potential abuse and the perception of abuse are real concerns, but we should not throw the baby out with the bathwater. Guidelines can demand added independent review for approval of especially risky studies. They can define more painstakingly what extreme circumstances would justify high risk to participants. They can create multiple independent levels of rigorous review for the quality of consent. They can set severe penalties for breaches. All that and more should at least be tried before we jump to banning all high-risk studies, including studies of tremendous humanitarian urgency. There is no need to undermine our ability to respond to pandemics and to other major diseases effectively by banning wholesale risky research that is otherwise justified on balance.

If anything, some potential for misuse exists if caps on risk to study participants are legislated. In recent discussions of challenge trials to fight coronavirus, several writers who had written on caps to risks in studies seemed to apply standards that were more demanding than their own writing on such caps had suggested. For example, instead of calculating net risks, as they had proposed (with some limitations), which would let background risks and indirect benefits of study participants reduce the overall risk being considered, when it was time to assess challenge trials, only direct risks explicitly entered the calculus (Rid, 2014; Palacios and Shah, 2019; Eyal, 2020; Shah *et al.*, 2020). Other writers seemed to

presume that a potential for death or serious injury in a study already rules it out (Deming *et al.*, 2020). In short, caps on acceptable risk can translate into more restrictions than they explicitly endorse.

Conclusion

This article both characterized the idea of caps on risk in medical and scientific research, and argued against their existence. Caps on risk to study participants, which forbid even studies of the utmost social value and urgency whenever their participants would face sufficiently high risks, appeal initially. But they break under the pressure of sustained analytical scrutiny. Threats to global public health could get even worse than the novel coronavirus. If we want to be able to fight these threats as effectively and expediently as we can, research ethicists should forego caps on risks in research—at least for well-defined extreme disaster circumstances, and with many checks in place. Study participants should instead be protected by other requirements, especially for informed consent, independent review and the permission to expose them only to net risks that are truly warranted by the social value of the study.

Funding

N.E.'s work was funded by NIAID (AI114617-01A1), NSF (2039320) and Open Philanthropy.

Conflict of Interest

None declared.

Notes

1. Normally, that process starts with exposure of a volunteer to a low dose, then a period in waiting to see if she gets infected, and does not develop severe disease, then another round with a higher dose and a similar wait period, then another with a still higher dose, and so forth. A faster but less safe alternative would be to up the dose by more in every round, or conduct some of the rounds synchronically.
2. Her list of cites is strikingly long: 'US Department of Health and Human Services (DHHS) 2009; Emanuel *et al.* 2000; Council for International Organizations of Medical Sciences (CIOMS) 2002; South African Medical Research Council (SAMRC) 2002; Council of Europe (CoE) 2005; Indian Council of Medical Research (ICMR) 2006; Ad

hoc group for the development of implementing guidelines for Directive 2001/20/EC (Ad hoc group) 2008; Schweizerische Akademie der Medizinischen Wissenschaften (SAMW) 2009; Schweizerische Eidgenossenschaft, Bundesamt für Gesundheit (BAG) 2011; World Medical Association (WMA) 2013' (Rid, :2014: 64).

3. Arrhenius and Rabinowicz (2014: 236). Note that they are discussing value superiority whereas we are discussing reason superiority.
4. But see the discussion below of the claim that medical studies always involve too much uncertainty to justify extremely high risk to study participants. In correspondence, Annette Rid suggested that doomsday scenarios that make conducting this or that study tremendously valuable may be popular with philosophers but remain sci-fi (email of 26 August 2018). Rid's position may thus be as follows: the social value of medical studies never surpasses some finite and modest level, and this is (part of) what keeps legitimate risks to study participants finite and modest.
5. An exception is a study for a countermeasure tackling a rare and mild disease that is anyhow about to be eradicated.
6. David Shaw writes of clinical trials, 'Even extreme risks can be acceptable provided they are described as such to potential participants' (Shaw, 2014: 1010). But Shaw is opposed to *any* oversight of the risk-benefit balance. Savulescu (2015: 103) proposes that in military research, soldier-participants 'could be granted leave from field duty to take part in similarly (and often extremely high) risky research'. But Savulescu anchors his position in the distinctive permission to expose soldiers to high risks in combat work, which does not obtain outside the special context of the military (on page 104, he mentions without argument that similar norms may be applicable to 'civilian' clinical trials). I am grateful to Annette Rid for both references. The nearest exception is (Steel, 2020) (2019), but even he seems to accept limits on study risk 'at the level of policy.
7. Compare Kolber (2020).
8. And the Walter Reed yellow fever study noted below illustrates that a high-risk, tremendous-value study can occur outside public health emergencies.
9. In studies that are less translational, uncertainty about translating into improved population health is higher. But the prospect of eventual translation into improved population health need not be lower as they usually have many more potential applications. An example is Walter Reed's Yellow Fever

study (see above), which led to scores of applications against mosquito-borne infections.

10. Email to the author, 20 June 2020.
11. Marcel Verweij suggested to me that caps on acceptable risk in medical studies may rest on skepticism that consent is ever fully free and informed, or at least that its fullness can be verified, even with strong measures to ensure high-quality consent. Inasmuch as that is the case, the propriety of caps on acceptable risk will depend on whether there exist caps on acceptable risk when consent is deficient, a question on which this article does not take a position.

Acknowledgements

The author is grateful to Annette Rid and Bastian Steuwer for detailed comments on an earlier draft. He also thanks Steven Joffe, Marc Lipsitch, Harisan Nasir, Robert Steel, and participants of the Post Research Ethics Analysis (PREA) conference (Ohio State University, 2019) for helpful suggestions.

Appendix

Caps on acceptable risk admits of many readings. While this article engages it on virtually all possible readings, distinguishing between several readings may advance the literature in the area:

1. *Ex-ante* vs. *ex-post* determinations: One natural reading of a cap on acceptable risk to study participants is that no participant can legitimately be exposed to personal risk of a certain level and above. But other readings are possible. In one alternative, the risk involved is the risk that, in the cohort of study participants, at least one severe adverse event would result from the study. On that reading, a study is illegitimate when the probability of bad harm to a *statistical* study participant is high enough (if only because the study has many participants, and the intervention is known to cause a rare severe complication). The language of the Nuremberg Code, 'No experiment should be conducted where there is an *a priori* reason to believe that death or disabling injury will occur' (International Military Tribunal, 1947: Article 5), is at least open to the latter reading. *Caps on acceptable risk* is compatible with either reading and the argument above attacks it under either.
2. *Aggregation-sensitive* vs. *aggregation-insensitive caps*: Usually, the cap on risk to participants is specified in terms of very severe individual harm. That a single participant is placed at high-enough risk is already considered a breach of the cap. But the cap may also be thought to remain sensitive to the number of

study participants at risk. At the extreme, it may forbid even a modest probability of modest harm when there are very many participants. As a case in point, can any large-scale health policy experiment of tremendous social importance justify denying an entire region fortification of drinking water, translating into modest impact on each but, in the aggregate, into much harm? Again, the argument above jeopardizes either answer to this question.

3. *Caps that weigh probability and harm equally vs. ones that place extra weight on one of them:* Usually, the risk levels specifying caps are understood, as *probability* times *harm*, where both are equal in weight. But there are alternatives. *Very* high probability of substantial harm may be thought worse than high probability of *very* substantial harm. For example, a study that carries a 100 per cent probability of costing a healthy person a tooth (i.e. it is certain to maim her) may be considered unacceptable even when a study carrying 0.1 per cent risk of death for that person might be considered acceptable. That could remain the case even if losing a tooth is deemed less than 1/1000 as bad as dying. The language of some of the quotations above emphasizes the researchers' level of confidence in their judgments that real harms will issue, so it may suggest such consideration. Conversely, a precautionary approach that rejects even a small chance of catastrophe may rule out even a small chance of true calamity. In this spirit, many object to coronavirus vaccine challenge trials that carry very low probability of a death of a study participant on the ground that still, a participant *may* die in those studies. Again, this difference is tangential to my argument against caps in all their variants.
4. *Caps specified in terms of net risks vs. ones specified in terms of risks, period:* 'Risks from a study' can designate either net risks, that is, the added risks to the participant minus the added benefits to her, or only the former. Several supporters of caps are explicit that they are discussing net risks. But the language of the Nuremberg Code, 'No experiment should be conducted where there is an *a priori* reason to believe that death or disabling injury will occur' (International Military Tribunal, 1947: Article 5), may decline to heed any benefits, e.g. in terms of preventing what would occur but for the experiment. Might the experiment hasten death but spare the participant of so much pain that good clinicians would have offered a similar intervention to a dying patient in her circumstances? Might it have caused the study participant's death but only after having extended her life [a possibility

shown in Lippert-Rasmussen (2001)]? Some philosophers place special emphasis on avoiding the introduction of risks, regardless of any benefits delivered (Shiffrin, 1999), so might reject a net-risk standard. Again, either reading is open to the critique advanced above. Let me end by demarcating the above distinctions from similar ones that apply instead to the social value of studies. In this vein, can a study's social value be considered high due to the aggregation of small benefits for many people (or, in an *ex-ante* version, to that of small improvements in prospects for many)? Some Scanlonian philosophers may prefer approaches that compare participants' (prospective) benefits to those of other individuals, in pairwise fashion, not in the aggregate. These philosophers might not consider the aggregation of many small benefits to many members of society sufficient justification for severe risk or harm to any individual participant. Likewise, the distinction between a standard risk calculation and one placing extra weight on probability can pertain to the assessment of social value. Many ethicists would be more willing to subject participants to high risk in a translational study that is one final step away from developing a countermeasure sure to make a modest dent in disease burden than in an exploratory study of bench science, which could generate helpful insights in many areas of medicine not one, with higher aggregate prospect of dramatic impact on human health, but without high probability of making any particular difference to it. For such ethicists, a bird in the hand is more important than many in the bush.

References

- Agrawal, M., and Emanuel, E. J. (2006). Phase I Oncology Research. In E. J., Emanuel, C., Grady, R. A., Crouch, R. K., Lie, F. G., Miller, and D., Wendler (eds), *The Oxford Textbook of Clinical Research Ethics*. Oxford: Oxford UP, pp. 356–366.
- Arrhenius, G., and Rabinowicz, W. (2014). Value Superiority. In I., Hirose, and J., Olson (eds), *Oxford Handbook of Value Theory*. Oxford: Oxford UP, pp. 225–248.
- Buchak, L. (2017). Why High-Risk, Non-Expected-Utility-Maximising Gambles Can Be Rational and Beneficial: The Case of HIV Cure Studies. *Journal of Medical Ethics*, **43**, 90–95.
- Cohen, J. (2016). Studies That Intentionally Infect People with Disease-Causing Bugs Are on the Rise, 18 May 2016. *Science*.
- Dawson, L., Earl, J., and Livezey, J. (2020). SARS-CoV-2 Human Challenge Trials: Too Risky, Too Soon. *Journal of Infectious Diseases*, **222**, 514–516.

- Deming, M. E., Michael, N. L., Robb, M., Cohen, M. S., and Neuzil, K. M. (2020). Accelerating Development of SARS-CoV-2 Vaccines—The Role for Controlled Human Infection Models. *New England Journal of Medicine*, 1 July 2020.
- Emanuel, E. J., Wendler, D., and Grady, C. (2000). What Makes Clinical Research Ethical? *Journal of the American Medical Association*, **283**, 2701–2711.
- Eyal, N. (2017). How to Keep High-Risk Studies Ethical: Classifying Candidate Solutions. *Journal of Medical Ethics*, **43**, 74–77.
- Eyal, N. (2020). Why Challenge Trials of SARS-CoV-2 Vaccines Could Be Ethical despite Risk of Severe Adverse Events. *Ethics & Human*, **42**, 24–34.
- Eyal, N., Lipsitch, M., and Smith, P. G. (2020). Human Challenge Studies to Accelerate Coronavirus Vaccine Licensure. *Journal of Infectious Diseases*, **221**, 1752–1756.
- Hay, M., Thomas, D. W., Craighead, J. L., Economides, C., and Rosenthal, J. (2014). Clinical Development Success Rates for Investigational Drugs. *Nature Biotechnology*, **32**, 40–51.
- HHS (2009). 45 CFR 46 (Human Subjects Research). USA.
- International Military Tribunal (1947). The Nuremberg Code. In *Trials of War Criminals before the Nuernberg Military Tribunals under Control Council Law No. 10. Nuernberg, October 1946-April, 1949*. Washington, U.S. Government Printing Office [1949–53].
- Jamrozik, E., and Selgelid, M. J. (2020). COVID-19 Human Challenge Studies: Ethical Issues. *The Lancet Infectious Diseases*, **20**, E198–E203.
- Jonas, H. (1969). Philosophical Reflections on Experimenting with Human Subjects. *Daedalus*, **98**, 219–247.
- Kagan, S. (1998). *Normative Ethics*. Boulder, Colorado: Westview.
- Kamm, F. (2001). *Inviolability*, Routledge Encyclopedia of Philosophy. available from: <https://www.rep.routledge.com/articles/thematic/inviolability/v-1>. [accessed 14 August 2020].
- Kolber, A. J. (2020). Why We (Probably) Must Deliberately Infect. *Journal of Law and the Biosciences*, **7**, Isaa024.
- Lippert-Rasmussen, K. (2001). Two Puzzles for Deontologists: Life-Prolonging Killings and the Moral Symmetry between Killing and Causing a Person to Be Unconscious. *The Journal of Ethics*, **5**, 385–410.
- London, A. J. (2009). Clinical Research in a Public Health Crisis: The Integrative Approach to Managing Uncertainty and Mitigating Conflict. *Seton Hall Law Review*, **39**, 1173–1202.
- Marston, H. D., Lurie, N., Borio, L. L., and Fauci, A. S. (2016). Considerations for Developing a Zika Virus Vaccine. *The New England Journal of Medicine*, **375**, 1209–1212.
- Miller, F. G., and Joffe, S. (2009). Limits to Research Risks. *Journal of Medical Ethics*, **35**, 445–449.
- Moore, M. (1989). Torture and the Balance of Evils. *Israel Law Review*, **23**, 280–344.
- Nagel, T. (1972). War and Massacre. *Philosophy and Public Affairs*, **1**, 123–144.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. New York: Basic Books.
- Palacios, R., and Shah, S. K. (2019). When Could Human Challenge Trials Be Deployed to Combat Emerging Infectious Diseases? Lessons from the Case of a Zika Virus Human Challenge Trial. *Trials*, **20**, 702.
- Parfit, D. (2011). *On What Matters*, Vol. I. New York: Oxford University Press.
- Plotkin, S. A., and Caplan, A. (2020). Extraordinary Diseases Require Extraordinary Solutions. *Vaccine*, **38**, 3987–3988.
- Pronker, E. S., Weenen, T. C., Commandeur, H., Claassen, E. H., and Osterhaus, A. D. (2013). Risk in Vaccine Research and Development Quantified. *PLoS One*, **8**, e57755.
- Resnik, D. B., and Sharp, R. R. (2006). Protecting Third Parties in Human Subjects Research. *IRB Ethics and Human Research*, **28**, 1–7.
- Rid, A. (2014). Setting Risk Thresholds in Biomedical Research: Lessons from the Debate about Minimal Risk. *Monash Bioethics Review*, **32**, 63–85.
- Rid, A., and Wendler, D. (2011). A Framework for Risk-Benefit Evaluations in Biomedical Research. *Kennedy Institute of Ethics Journal*, **21**, 141–179.
- Savulescu, J. (2015). Science Wars—How Much Risk Should Soldiers Be Exposed to in Military Experimentation? *Journal of Law and the Biosciences*, **2**, 99–104.
- Shah, S. K., Kimmelman, J., Lyerly, A. D., Lynch, H. F., McCutchan, F., Miller, F. G., Palacios, R., Pardo-Villamizar, C., and Zorrilla, C. (2017). *Ethical Considerations for Zika Virus Human Challenge Trials*. National Institute for Allergy and Infectious Diseases. Bethesda, MD: NIH.
- Shah, S. K., Miller, F. G., Darton, T. C., Duenas, D., Emerson, C., Lynch, H. F., Jamrozik, E., Jecker, N. S., Kamuya, D., Kapulu, M., Kimmelman, J., MacKay, D., Memoli, M. J., Murphy, S. C., Palacios, R., Richie, T. L., Roestenberg, M., Saxena, A., Saylor, K., Selgelid, M. J., Vaswani, V., and Rid, A. (2020). Ethics of Controlled Human Infection to Study COVID-19. *Science*, **36**, 832–834.

- Shaw, D. (2014). The Right to Participate in High-Risk Research. *Lancet*, **383**, 1009–1011.
- Shiffrin, S. V. (1999). Wrongful Life, Procreative Responsibility, and the Significance of Harm. *Legal Theory*, **5**, 117–148.
- Steel, R. (2020). Reconceptualising Risk–benefit Analyses: The Case of HIV Cure Research. *Journal of Medical Ethics*, **46**, 212–219. 10.1136/medethics-2019-105548
- Thomson, J. J. (1990).. Cambridge, MA: Harvard University Press
- US Centers for Disease Control (2018). *Remembering the 1918 Influenza Pandemic*. Atlanta, GA: CDC.
- Vannice, K. S., Cassetti, M. C., Eisinger, R. W., Hombach, J., Knezevic, I., Marston, H. D., Wilder-Smith, A., Cavaleri, M., and Krause, P. R. (2019). Demonstrating Vaccine Effectiveness during a Waning Epidemic: A WHO/NIH Meeting Report on Approaches to Development and Licensure of Zika Vaccine Candidates. *Vaccine*, **37**, 863–868.
- WHO (2018). Middle East Respiratory Syndrome Coronavirus (MERS-CoV). In *Factsheets*. Geneva: WHO.
- WHO Working Group for Guidance on Human Challenge Studies in COVID-19 (2020). *Key Criteria for the Ethical Acceptability of COVID-19 Human Challenge Studies*. Geneva: WHO, p. 20.
- Wikler, D. (2017). Must Research Benefit Human Subjects If It Is to Be Permissible? *Journal of Medical Ethics*, **43**, 114–117.
- Williams, B. (1973). A Critique of Utilitarianism. In B., Williams, and J. J. C., Smart (eds), *Utilitarianism—For and Against*. Cambridge: Cambridge University Press, pp. 77–150.