Why Restrict Medical Effective Altruism?

Travis Quigley

Quigley, Travis (2024). Why restrict medical effective altruism? Bioethics 38 (5):452-459. 10.1111/bioe.13279

Keywords: Challenge Trials, Effective Altruism, Research Ethics, Ethics of Risk

Abstract: In a challenge trial, research subjects are purposefully exposed to some pathogen in a controlled setting, in order to test the efficacy of a vaccine or other experimental treatment. This is an example of Medical Effective Altruism (MEA), where individuals volunteer to risk harms for the public good. Many bioethicists rejected challenge trials in the context of Covid-19 vaccine research on ethical grounds. After considering various grounds of this objection, I conclude that the crucial question is how much harm research subjects can permissibly risk. But we lack a satisfying way of making this judgment, that does not appeal simply to the intuitions of doctors or bioethicists. I consider one recent and structurally plausible approach to critically evaluating the harm question. Alex London defends a social consistency test for research risks: we should compare the risks undertaken by research subjects to relevantly similar risks which are accepted in other spheres of society. I argue there is no good reason not to consider volunteer military service as a relevant social comparison. This implies there is essentially no cap on acceptable risks on the social consistency rationale. In short, if soldiers can be heroes, why can't research volunteers?

In a challenge trial, research subjects are purposefully exposed to some pathogen in a controlled setting, in order to test the efficacy of a vaccine or other experimental treatment. Challenge trials became recently prominent in the Covid-19 pandemic, when some believed that they might meaningfully accelerate vaccine development.[1] Many bioethicists rejected challenge trials in that context on

[1] The organization 1Day Sooner gathered numerous volunteers — including myself — from around the world. See https://www.1daysooner.org/.

specifically ethical grounds, especially that the harms to subjects would be excessive.[2] Blocking challenge trials on patient-centered grounds, when both researchers and volunteer participants are ready to move forward, is an instance of paternalistic intervention to block what we can call *medical effective altruism* (MEA). There are other examples of MEA, such as regulations around voluntary organ donation,[3] but I mainly focus here on challenge trials as an example of MEA in a research context.

There is a certain tendency toward conservatism among bioethicists in such cases, which has been recently criticized by some philosophers.[4] Challenge trials are an accepted research practice in many contexts, but there was significant hesitation in the novel context of Covid-19, because the risks of challenge trials could not be precisely estimated, were presumed to be too high, or both. The crucial

---

[2] See, e.g., Dawson, V., Earl, J., & Livezay, J. (2020). Severe Acute Respiratory Syndrome Coronavirus 2 Human Challenge Trials: Too Risky, Too Soon. *The Journal of Infectious Diseases.* 222(3), 514-516. Shah, S., Miller, F., Darton, T., Duenas, D., Emerson, C., Lynch, H...Rid, A. (2020). Ethics of Controlled Human Infection to Address COVID-19. *Science.* 368(6493), 832-834. Solbakk, J., Bentzen, H., Holm, H., Heggestad, A., Hofmann, B., Alnaes, A...Bernabe, R. Back to WHAT? The Role of Research Ethics in Pandemic Times. *Medicine, Health Care, and Philosophy.* 1(3), 3-20. Kahn, J., Henry, L., Mastroianni, A., Chen, W., Macklin, R. For Now, It's Unethical to Use Human Challenge Studies for SARS-CoV-2 Vaccine Development. *Proceedings of the National Academy of Sciences.* 117(46), 28538-28542. For some additional contemporaneous reactions, see Branswell, H. (2020). Infect Volunteers with Covid-19 in the Name of Research? A Proposal Lays Bare a Minefield of Issues. *Stat News*. Retrieved from https://www.statnews.com/2020/05/01/infect-volunteers-with-covid-19-in-the-name-of-research-a-proposal-lays-bare-a-minefield-of-issues/

[3] For instance, see discussion of the "dead donor rule": Sade, R. & Miller, F. (2014). Consequences of the Dead Donor Rule. *The Annals of Thoracic Surgery.* 97(4), 1131-1131; Zimmerman, C., Baggett, N., Taylor, L., Buffington, A., Scalea, J. Fost, N...Schwarze, M. Family and Transplant Professionals' Views of Organ Recovery before Circulatory Death for Imminently Dying Patients. *American Journal of Transplantation: Official Journal of the American Society of Transplantation and the American Society of Transplant Surgeons*. 19(8), 2232-2240; and Mezrich, J. & Scalea, J. (2015). As They Lay Dying. *The Atlantic.* Retrieved from https://www.theatlantic.com/magazine/archive/2015/04/as-they-lay-dying/386273/

[4] e.g., Chappell, R. & Singer, P. (2020). Pandemic Ethics: The Case for Risky Research. *Research Ethics*. 16(3-4), 1-8. Chappell, R. (2022). Pandemic Ethics and Status Quo Risk. *Public Health Ethics*. 15(1), 64-73.

question for this paper is how much harm research subjects can permissibly risk.[5] There have been multiple recent arguments criticizing the idea that there should be any absolute cap on permissible harms in emergency circumstances where the benefits of risky research might be enormous.[6] I wish to raise a set of questions that is broader in two ways. First, the issue of MEA is not solely relevant in emergency circumstances. What is the underlying ethical standard in normal times, and what are its grounds? Second, while defeating the idea of an absolute cap on research risks is important (and by no means uncontroversial), MEA can occur at low and medium stakes as well as high. What are the appropriate standards for lower-stakes challenge trials, or for altruistic organ donation, which may be live-saving but is not a matter of social crisis?

There is a further methodological aim of this paper. Arguments against challenge trials tend to combine many different objections. This is an understandable argumentative strategy if one's goal is to thoroughly rebut challenge trial advocates. But it blurs the underlying issues. I attempt to isolate the excessive harm question, which I believe is the central bioethical concern, from various other objections. Section 1 commences this survey, after giving a preliminary argument that MEA should be *prima facie* acceptable because it is autonomous, beneficent action, the restriction of which amounts to paternalism. Section 2 develops a framework for weighing acceptable risks. Alex London proposes a social consistency test for research risks.[7] Roughly, we should look to socially accepted risky practices in other domains to formulate standards for the medical domain. I argue that London does not carry the argument far enough: once we make natural comparisons to military or emergency service, there is no basis for paternalistic restriction on MEA. If soldiers can be heroes, why can't research volunteers?

---

[5] This is most naturally represented as a cap on the overall expected harms of a given act. Of course, almost all medical research, as with anything in life, has *some* minimal risk of severe outcomes including death. So risk assessments must generally be done based on expected harms, rather than on the absolute worst that could happen. I take for granted that this is what everyone discussing the ethics of risk has in mind.

[6] See Eyal, N. (2020). Is There an Ethical Upper Limit of Risks to Study Participants? *Public Health Ethics.* 13(2), 143-156.

[7] London, A. (2022). *For the Common Good: Philosophical Foundations of Research Ethics.* Oxford: Oxford University Press.

1. Preliminary Arguments

1.1 Altruism and Paternalism

The frequency of bioethical objections to medical effective altruism is puzzling on the face of things. Instances of MEA are, by stipulation, freely chosen and beneficent acts. The standard reasons for restricting free acts are to prevent harm to others, to prevent non-voluntary harms to self, or to prevent fully voluntary harms to self. The distinction between voluntary and non-voluntary marks the difference between "hard" and "soft" paternalism, where soft paternalism acts to prevent only non-voluntary harms to self.[8] Acts of MEA clearly don't harm others, so the only relevant justifications are paternalistic.

The idea of soft paternalism is that it seems acceptable to restrict actions that, while superficially free, do not actually accord with an agent's best interests as they understand them. Standard cases are ignorant or mistaken (by their own lights) choices. Acts of MEA would need to meet at least the normal standard of informed consent, the point of which is to ensure voluntary action.

The idea of hard paternalism is that some choices are harmful enough that they should be restricted even when fully voluntary. Standard cases are risky choices where an agent's preference clearly seems to outweigh any potential benefit, such as an insistence on solo skydiving without training. But MEA requires that an overall cost/benefit calculation be satisfied: the point of the research is that it is *effective* in overall public health terms. It would not make sense for a challenge trial, or any other research, to go forward if there were some strictly superior public health alternative. So the ordinary hard paternalist rationale doesn't apply. Some other possible cases of hard paternalism involve actions that seem repugnant or depraved, even if we acknowledge that they are free and beneficial: perhaps a practice of voluntary (mutual?) cannibalism, managed prudently to mitigate health risks and (by stipulation) highly enjoyable to the participants, would qualify. Whatever the status of such cases, clearly altruistic action isn't repugnant in this sense. Therefore, while the bioethical presumption

[8] See Feinberg, J. *The Moral Limits of the Criminal Law 3: Harm to Self.* (1989). New York: Oxford University Press, p. 12.

seems to be against MEA, this gets things backwards: free and beneficent action is presumptively favored.

What would this imply, as a *prima facie* position on risky research? We should draw a distinction between reasonable and unreasonable altruistic action. Altruism is reasonable when the expected benefits, considered agent-neutrally, outweigh the costs. Sacrificing yourself in a cause that does not impartially justify the costs is normally a moral mistake.[9] While it is not obvious when such a mistake warrants intervention, in cases of MEA, there is no reason for third parties to be involved in acts that are not impartially justified at all — a challenge trial that does not meet a cost/benefit test is simply not a viable candidate.[10] Given that MEA is impartially justified — perhaps even if only barely — and informed consent is obtained (see 1.4 below), the *prima facie* position, given a presumption against paternalism, is that MEA should be permissible. I now survey some objections which attempt to unseat this *prima facie* position — although many advocates argue as if the burden of proof is opposite of what I have suggested.

1.2 Research Benefit and Design

Challenge trials would obviously not be justified if the expected public health benefits did not impartially justify the risks. As mentioned already, I will generally stipulate that the scope of this argument concerns only research that would be distinctively socially beneficial. But two points warrant brief discussion here.

The first point is that, while MEA is impartially justified by hypothesis, one might question whether high-risk challenge trials are likely to ever meet this condition. This would not affect the structure of my argument, but would reduce the relevance of my chosen example. Here are two

---

[9] For discussion of the distinction between "reasonable imprudence" and "reckless imprudence," see Jansen, L. & Wall, S. (2018). Reconsidering Paternalism in Clinical Research. *Bioethics*. 32(1), 50-58.

[10] This raises a certain complication. I have said that blocking MEA is paternalistic. But all MEA requires cooperation of third parties, so withdrawing that cooperation is not, strictly speaking, blocking an altruist from choosing an option they have — it is failing to provide them with an option they want. However, MEA research by hypothesis provides an impartial public health benefit, and medical providers have a presumptive reason to promote public health, so the only patient-centered reasons not to provide an MEA option are based on paternalistic prevention of altruistic self-sacrifice.

observations. The first is that challenge trials are an accepted practice in research that has a lower risk of patient harm.[11] There is no clear reason why the prospect of research benefit would drop away as risks of harms to subjects rise. Indeed, it is not hard to imagine scenarios where challenge trials might be especially important in high risk cases. The primary rationale for challenge trial advocacy regarding Covid-19 was largely superseded, because the rapid spread of the virus meant that conventional trials (which are based on virus exposure "in the wild") proceeded quickly.[12] But it is plausible that a more ideal pandemic response would begin with short- or medium-term lockdown tactics which seriously mitigate spread, and conduct challenge trials in *that* context, because the very success of lockdowns would slow down conventional vaccine trials. It is important for future pandemic preparedness to consider the ethical permissibility of this strategy.

The second point is about the appropriate sequence of ethical reasoning about challenge trials. Bioethical arguments for restricting MEA, as I've noted, tend to evaluate several different moral factors at once. This is understandable, but can lead to odd results. An uncontroversial ethical provision is that risks to research subjects should be minimized so far as possible: challenge trial volunteers, for instance, should have ready access to the best available treatments. A common proposal is that subjects should be limited to the young and healthy, because they are likely to face the lowest risks from infection.[13] But this can be turned against challenge trial proposals: Kahn et al argues that the research benefits of challenge trials are overstated precisely because their results, based on young healthy people, would not clearly generalize to more vulnerable populations.[14] This gets the ethical order of operations wrong. We

[11] See section 2 below, and Miller, F. & Grady, C. (2001). The Ethical Challenge of Infection-Inducing Challenge Experiments. *Clinical Infectious Diseases.* 33(7), 1028-1033.

[12] There are some possible secondary use cases: for instance, the UK altered the dosing schedule of vaccines to prioritize "first doses first" by increasing the interval between first and second doses. See Imai, N., Rawson, T., Knock, E., Sonabend, R., Elmaci, Y., Perez-Guzman, P...Cori, A. Quantifying the Effect of Delaying the Second COVID-19 Vaccine Dose in England. *The Lancet Public Health.* 8(3), 174-183.

[13] For a discussion that compares the infection fatality rate of Covid-19 *for young people* to rates of complications from organ donation, see Jayaram, A., Sparks, J. & Callies, D. (2022). Justifying the Risks of COVID-19 Challenge Trials: The Analogy with Organ Donation. *Bioethics*. 36(1), 100-106.

[14] Kahn et al, *op. cit* note 3.

must first identify the research profile that has the best cost/benefit profile, factoring in both research benefits and harms to subjects. Only then can we ask whether the proposal violates ethical side constraints, in particular by exposing subjects to unacceptable harms.[15] Of course, justifying the inclusion of more vulnerable populations, at the research design stage, would require commensurately higher expected benefits. But we should not presume a certain risk minimization practice at the research design phase that undermines the research benefits of the proposal. This conflates the question of research design and the question of side constraints on acceptable harms to subjects.

1.3 Priority of the Individual

Solbakk et al make a series of objections to challenge trials.[16] One idea is that risky challenge trials are ruled out because of the principle that "the interests of the individual research participant is paramount in research ethics...the interests of society has to yield to the interests of the participant."[17] This principle is found, the authors point out, in various "normative documents," such as the Nuremberg and Helsinki Declarations. The objection is akin to Rawls's famous "separateness of persons" objection to utilitarianism, which holds against weighing the interests of an aggregate of people against the interests of an individual.[18]

The nature of this claim is not entirely clear, especially when we are considering voluntary action. We should not rule out that an individual finds altruistic sacrifice to be in their own interest: this is to rule preemptively in favor of a narrow self-interest conception of rationality. In other cases, such as military service, we recognize that an individual may choose to sacrifice themselves for the benefit of others.

---

[15] Kahn et al, ibid, suggest an additional side constraint, about the acceptability of enrolling subjects from groups that are both vulnerable and historically disadvantaged, in particular Black Americans. This question is beyond the scope of this paper. But the rationale given is a particular application of the institutional trust argument, discussed below, and my misgivings about that argument — that it is highly speculative and not properly an issue of bioethical expertise— apply.

[16] Solbakk et al, *op cit* note 3.

[17] Ibid: 14.

[18] Rawls, J. *A Theory of Justice*. 1971. Cambridge, Mass: Belknap.

The authority of "normative documents" is, of course, philosophically questionable. In particular, many of the restrictions on medical research are motivated by research atrocities, such as those conducted by Nazis, which were obviously *not* licensed by informed consent. It is at the least not obvious that such standards have the last word to veto consensual altruistic action.

1.4 Informed Consent

A possible rejoinder is that informed consent, in the relevant cases, is impossible, and therefore consent cannot be deployed as I just did. Some argue that informed consent is impossible in cases like Covid-19 challenge trials because the risks cannot be adequately known by researchers or patients.[19] This implies a dubious idea of informed consent. We take actions all the time with the understanding that there are some "known unknown" risks involved. Knowing the precise character of the various harms I might suffer is an impossible standard of informed consent — taking an antibiotic may lead to catching a fever due to my reduced immune function, which could lead to nearly anything.[20] The reason this objection is compelling must be due to the severity of the risks, not the uncertainty about their nature. It is true that, in the case of most conventional medical procedures, there are statistical generalizations available about the severity of the risk profile that are not available for experimental work. But it is perfectly possible to be informed *that* the degree of likely harm is uncertain. If participants consent to an unknown degree of risk, the operative ethical issue must be the (possible) degree of expected harm itself, not the possibility of consent.

---

[19] Kahn et al, Solbakk et al, *op. cit.* note 3.

[20] Cases can be multiplied. Every time I get into a car, I risk essentially any possible injury, including serious brain damage and attendant (unimaginable) changes to cognitive function, etc. The first astronauts in space faced unknowable risks of surprising changes to the human body in those conditions. Surely these choices are consensual. Another set of cases in the medical realm is participation in experimental treatments: these are given special dispensation as "last resorts," but nonetheless involve informed consent to radically unknown possibilities.

There is a second set of informed consent worries, concerning the possibility of a "therapeutic misconception" or an analogous "preventive misconception."[21] The idea is that patients might participate in research trials because they believe they will receive some health benefit — a therapeutic benefit for some ailment they have, or the prevention of illness in the vaccine case — even when researchers are clear that there is no expected benefit. While the therapeutic misconception may be troubling in some cases, it would be odd to categorically block challenge trials on this basis. Perhaps we might demand a particularly clear and thorough type of informed consent to serve as a safeguard here, but surely there is *some* informed consent protocol that would suffice. The only reason I see to doubt this is a strong presumption that, if someone volunteers for MEA, then they *must* have some misunderstanding. But this is simply to presume that altruism can never be autonomously chosen, or at least that there is no evidence that could ever suffice to demonstrate truly voluntary altruism.

1.5 Institutional Trust

Another common worry about challenge trials is that patients becoming ill and dying in clinical research would undermine trust in public health institutions, and perhaps in particular reduce vaccine take-up rates, which are already a serious problem.[22]

This is a plausible concern, but it should be separated from the other bioethical objections to challenge trials, or other MEA, for two reasons. First, it is a speculative concern at the level of mass communication and public behavior — more in the domain of sociology than normative ethics. (It seems at least plausible that faster vaccine development might, with suitable transparency about the nature of the MEA sacrifice being made by volunteers, *increase* institutional trust.) Second, the harms that might occur in this way do not lie in the domain in which bioethics has the final say. Imagine that we were confident that some challenge trial would be highly valuable research, *and* that conducting the trial would lead to decreased public trust in public health institutions to a certain degree. It is certainly

---

[21] On the therapeutic misconception generally, see Miller, F. & Joffe, S. (2006). Evaluating the Therapeutic Misconception. *Kennedy Institute of Ethics Journal.* 16(4), 353-366. Kahn et al, *op. cit.* note 3, raise the issue of the preventive misconception for the case of vaccine challenge trials.

[22] For instance, Shah et al, Kahn et al, *op. cit.* note 3.

intelligible that a small reduction in public trust might be justified by a certain research benefit. Who should evaluate this tradeoff? If we grant, for this discussion, that the challenge trial in question is *otherwise* ethically permissible — setting aside the concerns about excessive harm and so forth — then the remaining deliberation is, it seems to me, a matter for policymakers and politicians. Or perhaps it is a question of normative political philosophy. In any case, the problem of institutional trust is not distinctively bioethical, even when the institutions in question happen to be public health institutions. (Obviously, analogous questions could arise with respect to military action, or banking regulation, or whatever.) So I believe the institutional trust question, while worthy of consideration and study, should be put to one side: we should first establish when challenge trials can be *otherwise* ethically acceptable, and leave the institutional trust question (alongside other practical questions, such as simple financial cost) to a later stage of moral and political inquiry.

2. The Harm Debate

The discussion so far has been preliminary. It is clear that the considerations discussed so far do not exhaust the field. Indeed, when we isolate the various objections, it seems clear that concerns about informed consent, trust, and so forth, draw much of their strength from the fact that challenge trials (and MEA in general) involve serious harms (or risks of harms) to research subjects.

Harm objections have at least two sub-categories, one that focuses more on risk of death and rescue therapies, and one more focused on serious discomfort or residual illness. These are sometimes presented as distinct, but seem to share the same core idea.[23] A rescue therapy would presumably prevent certain kinds of severe harm, so the relevance of such a therapy is simply that expected harms would be reliably prevented, or at least reliably (although likely not perfectly) capped in their severity. A further possible distinction is between "harm" and "discomfort," where discomfort is not associated with serious threats to health beyond the proximate experience of symptoms.[24] Since I will be arguing in favor of permitting risks of serious harms, including death, I will set aside these complications and simply speak in terms of harm.

---

[23] e.g., Solbakk (*op. cit.* note 3: p. 14) mentions rescue therapies separately from the discussion of harm.

[24] See Miller, F. & Grady, C., *op cit* note 11: p. 1030.

There are two important arguments in favor of challenge trials which I should distinguish from my own approach. The first is from Eyal, who argues that we should reject a hard upper limit on acceptable research risks, especially in emergency circumstances, because we should reject "moral absolutism."[25] The basic claim, with which I agree, is that a broad consensus in moral philosophy rejects the idea of morally absolute rules, which can *never* be violated. It is at least conceivable that, in emergency circumstances, a challenge trial might be extraordinarily valuable. If we reject moral absolutism, and accept the conceptual possibility that a challenge trial might be extremely valuable, then we should reject the idea of a hard upper limit on research risks.

This argument, even if accepted, is limited in certain ways. Most obviously, it pertains only when the stakes are very high. This prompts in response the "research benefit" objection which denies that the real benefits of any challenge trial could be high enough. Further, the question of acceptable risks and harms is not relevant only at extremely high stakes. There are many open questions remaining after the denial of moral absolutism. The social comparison strategy to follow will apply along the full spectrum.

Further, practically speaking, this argument seems to rely on a distinction between ordinary rules for non-emergency circumstances, and exceptions which can be made in emergencies. Solbakk et al argue that we should not be in the business of "pandemic exceptionalism," which relaxes normal ethical standards in extreme times.[26] One worry is that exceptionalism excessively empowers those in leadership roles, who may be under significant pressure to act in ways that are not ultimately defensible. I agree that we should not engage in pandemic exceptionalism, and should instead evaluate our background normative ideas in this area.

A second important argument, which is a close neighbor of my own argument here, is Jayaram et al's argument that challenge trials may be acceptable by comparison with altruistic (non-directed) organ donation.[27] This argument starts with one accepted kind of medical effective altruism and argues

---

[25] Eyal, *op. cit*. note 6: p. 147. Also see Hausman, D. (2021). Challenge Trials: What Are the Ethical Problems? *The Journal of Medicine and Philosophy*. 46(1), 137-145.

[26] Solbakk et all, *op. cit.* note 3.

[27] Jayaram et al, *op. cit.* note 13.

by consistency reasoning that at least some risky challenge trials may be acceptable. The dialectical appeal of this argument also represents a limitation: it relies on the internal standards of current bioethics, and therefore cannot shed light on what those standards *should* be. It is indeed telling that some MEA is accepted, but it is not obvious why, say, liver donation should serve as an upper limit for acceptable harms. The forthcoming argument is also comparative, but widens the scope to accepted practices within society as a whole, not just existing bioethics.

A broad social consistency approach looks to the ethical norms implicit in society overall, rather than just in bioethics. This provides a standpoint for critical evaluation of bioethical questions, without appealing to some controversial moral doctrine only accepted by a few (as would be the case if we sought to directly appeal to a doctrine in moral philosophy such as utilitarianism or Kantian ethics). In particular, the idea here is to look at other risky behaviors which are widely accepted in society, in order to make comparisons about medical risks.

The idea of a critical standpoint for bioethics is important here, because claims about acceptable harm are often vague and deferential to bioethical authority — either historical "normative documents," as we saw above, or expert committees[28] or institutional review boards, likely to be composed of members steeped in conventional bioethical precepts.[29] While ultimately any ethical approach will need to be institutionalized in some form, which will involve expert consultation, the specificity and principled basis provided for oversight is crucial. For instance, in Miller and Grady's important discussion, they conclude with a "spectrum" of examples: mild discomfort from challenge trials is clearly acceptable; trials where "treatment is nonexistent or ineffective [or] symptoms are intolerable" are clearly unacceptable; and a controversial range lies between.[30] But, while they consider a range of relevant issues, the key terms of the spectrum are vague — when is a symptom intolerable? is there a sharp distinction between discomfort and harm in practice? what constitutes a "small risk" of

---

[28] e.g., Shah et al, *op. cit.* note 3.

[29] e.g., an intuitive weighing of Beauchamp and Childress's famous four principles of bioethics — of which three (autonomy, beneficence, and justice) seem plausibly to weigh in favor of MEA, while only non-maleficence clearly weighs against. See Beauchamp, T. & Childress, J. (2019). *Principles of Biomedical Ethics*. New York: Oxford University Press.

[30] Miller, F. & Grady, C., *op. cit.* note 11: p. 1032.

death? — and the very starting point, that serious risks of harm are unacceptable, is based simply on "current knowledge," which seemingly refers only to a kind of bioethical common sense. This sort of framework reflects, rather than critically evaluates, bioethical doctrine.[31] This provides important motivation for an independent social consistency approach, which looks to the broader norms of society rather than the internal norms of bioethics.

London has recently developed this kind of argument.[32] This is a step in the right direction, but, as we will see presently, London does not go nearly far enough. His basic approach for evaluating research risks is to draw comparisons with "structurally similar" activities which are broadly accepted by society.[33] His chosen examples are volunteer (unpaid) firefighter or paramedic services, because "the volunteer nature of these activities combined with their orientation to serving the public interest represent important structural similarities to the research enterprise."[34] This is a fairly restrictive comparison, which seems like it would clearly rule out cases such as Covid-19 challenge trials.[35]

The results of the social consistency argument clearly rest on which activities are deemed relevantly similar comparisons — as London says, this is "an inherently normative or evaluative process."[36] I want to begin in a different place than London, and look to the riskiest choices which are socially permitted (and indeed lauded): the obvious candidate is voluntary (but professional) military service. The basic question is, given that individuals are allowed to join the military even in times of active war, why shouldn't they be allowed to volunteer for just about any act of MEA that we can imagine? Given the obvious high risks of injury and death in war, this would permit (e.g.) challenge trials even for *very* dangerous pathogens, if the overall altruistic case can be made. (And the altruistic case for challenge trials can get *stronger* the more deadly the virus.) This makes a very straightforward

---

[31] This discussion is cited approvingly by Solbakk et al, *op. cit.* note 3: p. 11.

[32] London, *op. cit.* note 7.

[33] Ibid: 260.

[34] Ibid.

[35] Indeed, he seems to think that even volunteer firefighters may take on too many risks to be fully appropriate comparisons. Ibid.

[36] Ibid.

social consistency argument against anything close to the conventional approach to voluntary risks. It is notable that military service is not merely acceptable, but is widely treated as laudable and worthy of veneration. And this is true, indeed *especially* true, when the risks of service are especially high (and seem especially worthwhile) — those who fought in Normandy on D-Day are regarded as American heroes.

To buttress the argument, further consider various risks taken in pursuit of entertainment: skydiving, downhill skiing, free solo rock climbing, and so on. Such thrill-seekers, while not lauded as heroes like soldiers, are permitted to do as they wish in a wide (though not unlimited) domain. This comparison, as well, seems like it would clearly license most cases of MEA — seemingly more than London is willing to countenance, since he rejects the comparison. The rest of this section is dedicated to discussing potential disanalogies between military service and thrill-seeking as compared to MEA. But I'll argue that the obvious disanalogies imply that *more* risk-taking should be permitted in cases of MEA than in the comparisons, not less. Finally, recall that our goal is to critically evaluate bioethical norms: so we should attempt to set aside, for these purposes, received intuitions about the role obligations of physicians (to do no harm) and so forth, lest we spoil the purpose of the exercise.

There is an important methodological point here. Social comparison arguments are a form of analogical reasoning — similar, as Jayaram et al note, to casuistry.[37] There is a tendency, in this process, to rely on comparisons that are *strictly similar*. This is exemplified by London's focus on structural similarity: the reason that volunteer first responders are chosen as an example is that they are unpaid — he does not envision significant compensation for challenge trials, which I discuss below — and focused on the public good. But he worries about "the most dangerous" risks of even volunteer first responders as a comparison, because of a "principal-agent" problem which is present in challenge trials but not (volunteer) emergency service. Thrill-seeking activities are set aside as comparisons because participants desire the thrill specifically, where challenge trial participation is not (one imagines)

[37] Jayaram et al, *op. cit.* note 13: p. 101, n. 7.

thrilling.[38] In a similar vein, Jayaram et al worry that the benefits of altruistic organ donation — their chosen comparison — are relatively certain, while the benefit of any one challenge trial is uncertain.[39]

In all of these cases, the presumption is that we should rely on *strict similarity*, where any descriptive difference between a comparator is evidence against it. But this is overly simplistic casuistry: we should instead focus on *relevant similarity*. In social comparisons, only some differences are relevant — for instance, we would need an explanation for why the certainty or uncertainty of the benefit matters at all, given that all other ethical parameters are satisfied. Further, some differences between cases can be just as telling as their similarities. I argue in the next paragraphs that the differences between military service and thrill-seeking are of a kind that make challenge trials appear, by comparison, *even more clearly acceptable*, once we evaluate the nature of the difference. Here is an illustration of the two kinds of similarity. I am trying to persuade someone to take a Covid-19 vaccine, and, to persuade them, I point out that they have taken many other vaccines. They point out that this vaccine is a new mRNA technology, and they are worried that it will misfire and give them Covid. Under strict similarity, this is a fair point. But, in actuality, I should observe that the mRNA vaccine — unlike previous vaccines — contains no viral material at all, so it is impossible for it to give them Covid-19 (even if it may have other side effects). If *that* is the basis of their fear, the disanalogy should make them *less* afraid of the mRNA vaccine.

Returning to the case at hand, we saw that London noted that the key structural features of volunteer firefighters is their volunteer (unpaid) nature and their orientation to the public good. Let me consider the military comparison first, which shares the public good motivation but is a paid profession. Using relevant similarity we should ask *why* the fact that soldiers are paid should render them an inapt comparison for MEA?

---

[38] London, *op. cit.* note 7: p. 260. The idea of the principal-agent problem is that the researcher is in a position of authority with respect to the subject. But, as above, this seems addressed by an adequate standard of informed consent.

[39] Jayaram et al, *op. cit* note 13: p. 101-102. One problem with Jayaram et al's comparison is that they ground the similarity, in part, on the shared "medical context" — this implies that it is *impossible* to use comparisons that provide a social perspective outside of the medical domain, if the shared medical context is itself treated as an analogical requirement.

There is a general bioethical concern about the role of financial compensation for medical risks or harms. Some kinds of compensation may be 'offers you can't refuse,' which might impinge on autonomy or on the status of informed consent. These would be deemed "undue inducement." I don't need to evaluate the merits of undue inducement here.[40] The critical point is that, even if we take concern about undue inducement from granted such that MEA cannot be significantly compensated, that doesn't mean that we shouldn't make comparisons to compensated activities in our social consistency framework. The implication is quite the opposite. If the risks of well-paid military enlistment are socially acceptable, then similar *un*paid risks should be *more* clearly acceptable, because we have stipulated away one site of ethical concern. In other words, in cases of MEA we know that financial motivation isn't doing the work, so the risks are chosen freely rather than out of need.

One could reverse the concern, and hold that it is *only* acceptable to take on significant risks when they are duly compensated, perhaps out of a concern for fairness — so, we actually *should* give significant compensation to challenge trial volunteers. But it would be odd to at once hold that view *and* a strong view prohibiting undue inducement. That would be to say that risks are only acceptable given an agency-undermining condition. If fairness requires compensation, but compensation undermines agency, then risk-taking is simply ruled out across the board, at least for the poor. This is a strange result, and obviously not one licensed by a social consistency argument. The underlying issue here is that undue inducement is a distinctively bioethical concern: people are not in practice overly concerned about military enlistees being unduly induced, despite the fact that some recruits are poor and have few good alternative opportunities. The worry about fairness may have more bite if we wish to generally reward socially esteemed behavior. But the impulse to reward is contrary to the impulse to avoid giving incentives: only an incoherent combination of the two contrary impulses would justify ignoring the military comparison.

The more general concern here is about the right motives for risky choices. As I noted, London also sets aside comparisons to thrill-seeking activities. Again, choosing to set aside the comparison

---

[40] I myself am not so sure about those merits. For critical discussion, see Wertheimer, A. (2011). *Rethinking the Ethics of Clinical Research*. New York: Oxford University Press.

rather than reasoning through the relevant differences leads to an unwarranted conservative conclusion. If risks taken merely to seek thrills are permissible, why would a similar risk, taken for altruistic reasons rather than amusement, be impermissible? The opposite position would be more coherent: perhaps thrill-seeking is a bad reason to take risks, and we should act paternalistically to restrict such behaviors. If that were so, we might yet hold that risk-taking for good motives (such as MEA or some military service) is permissible. Or, more radically, we might hold that risk-taking past a certain point is always unacceptable. But it is bizarre to hold that a given level of risk is acceptable with bad motives, but unacceptable with good motives. So thrill-seeking activities should not be set aside: if anything, they should be used as a minimal *baseline* for acceptable MEA risks. This is structurally similar to the point about monetary compensation: if we being with a socially accepted risky behavior and then eliminate a concern about bad motives, we should then be even more accepting of risky choices.

London's choice of comparisons, then, acts as a way of stacking the deck. Perhaps he would simply accept this point. He writes that our theoretical goal is to "find reasonable criteria of similarity…This process itself may require careful adjustments in the criteria of similarity…This is therefore an inherently normative or evaluative process."[41] This suggests that the judgment of structural similarity is being heavily shaped by a process of "reflective equilibrium," in which the outcomes of the argument are dictated by their intuitive appeal in addition to theoretical considerations. This is both understandable and common, but we should be wary: if the "adjustments" in the relevant comparisons are driven by underlying intuitions common in bioethics, then we will have lost the independent critical standpoint that the social consistency approach promised to deliver.

I believe a straightforward approach to the social consistency argument yields the conclusion that MEA risks are comparable to risks such as military service. That is a radical conclusion relative to the bioethical consensus, because it implies essentially no hard cap on what risks are acceptable. But the parameters of the conclusion are narrow: it applies only to freely chosen and beneficial actions. In that

[41] London, *op. cit.* note 7: p. 262.

sense the argument is very simple: volunteering to take drastic research risks may be a kind of heroic sacrifice of the kind soldiers make by fighting in a just war.

There is something further that is notable regarding the broader social comparisons to military service and thrill-seeking activities. It is not the case that risks in the line of military duty are generally unacceptable in normal circumstances, but that certain emergency conditions render risks acceptable. This may be overlooked, because it is easiest to illustrate acceptable military interventions with extreme cases such as fighting against Nazi Germany in World War II. But lower stakes risks can also be justified — consider humanitarian interventions, such as rescue work in natural disasters. These types of operations are not ordinarily as risky as a shooting war, but there are no particular threshold effects. Instead, there appears to be a proportionality standard at all levels — notably including paid (rather than volunteer) first responders like police officers and firefighters. I focus on this kind of example because such jobs clearly seem justified, and rightly socially respected, and some significant degree of risk seems inevitable. I believe consideration of the broader range of social comparisons returns us to the *prima facie* position with which I began: if a given risk is a genuine case of altruism, and is accepted freely, there is no sound basis for paternalistic restriction on grounds of excessive harm.

3. Conclusion

I argued that MEA has *prima facie* presumption that it should be permitted, based in the idea that free, beneficent action should not be paternalistically restricted. I then considered a range of objections. Let me summarize the state of play as I see it. The objections from the priority of the individual and informed consent seem to me misguided: a suitable standard of informed consent must be possible for MEA unless we take up a radical kind of skepticism about consent, and voluntary MEA does not involve a troubling form of moral aggregation (precisely due to being autonomously chosen). The objections from research benefit and institutional trust are plausible, but are not reasons to presumptively dismiss challenge trials or other MEA on the basis of excessive risks to subjects. Whether research benefits outweigh the risks to the health of subjects is ultimately an empirical question (albeit a potentially complex one). Whether research benefits (more accurately, the *net* benefits) would justify

any possible reduction in public trust is a question to be left for another day, and perhaps answered by policymakers or political philosophers.

The core bioethical objection, then, concerns acceptable harms to MEA volunteers. Existing claims on this score are often vague, or consist largely in appeals to authority, or (usually) both. I argued for employing a social consistency test, and in particular a broad social consistency test that includes activities outside the medical domain. The most developed social consistency test in the literature, from London, falls considerably short of the full potential of the strategy, because it uses an impoverished form of analogical reasoning — *strict* similarity rather than *relevant* similarity. Once we compare socially accepted comparisons, and take into account their relevant differences, there is no social consistency argument against justified, autonomous MEA. This argument, importantly, does not turn on emergency conditions or extraordinarily high stakes. It is, instead, a general independent argument that MEA is broadly acceptable by bioethical lights.