How Research Harms

October 17, 2023

Abstract

Scientific research harms us. We largely lack, however, a clear understanding of what those harms are, and under what contexts they instantiate. This is due in large part to a two-fold bias in science and policy which, first, takes most scientific harms to be the result of applying scientific research to the real world, rather than the research itself. Second, where we do consider the harms done at the research stage we focus heavily on harms to test subjects in the course of gathering data. While some philosophers have tackled instances of research harm outside of this context, there isn't any broad work that gathers and contextualises these projects under a clear structure. Here I create a preliminary classification system for research harms, ranging from physical harms to human test-subjects, to psychological, social, and moral hazards, aimed at clarifying how these harms instantiate and how to identify them for use in science and policy work.

1 Introduction

Science can harm us. Science has harmed us. Science is going to harm us again.

This isn't a new idea. Science, as a socially embedded, and influenced, field of study is part of a broad network of ideas and technologies that have

profoundly affected the world around us. This recognition - that science needs to think carefully about its impacts, both beneficial and detrimental - is part of an ongoing conversation in philosophy, the social sciences, and the sciences more broadly, on how to minimise harms while retaining advantages (Douglas, 2021; Nelkin, 1979; Pamuk, 2021). Whether it involves regulating the extractive practices of archaeology, limiting the ecological impact of building large-scale physics experiments, or reducing the risks to patients in medical trials, science is has undergone a significant shift in the way it thinks of itself in relation to the rest of society.

There are, however, two significant gaps in this conversation - first almost all discussion of the harms of science occur in the context of application. That is, we most commonly discuss, and attempt to control, scientific harms that occur when research is moved outside the lab and into practice - we ban the building of military AI, rather than AI research; geo-engineering deployment, but not modelling; gene-drives in viable human embryos, but not in labs. Indeed, there is a pervasive view that knowledge, and thus research, is value neutral, and generally ought not to be restricted.

Second, this, in turn, means that where we do consider minimising the harms that happen at the research level, we most often focus on the direct harms of experimentation - namely the harms done to animals or humans, instead of the broader harms that seemingly 'neutral' research can cause. Such harms do exist, however, and are are multi-modal - ranging from a lack of recognition of the immediate harms being done by research, to broader failures to regulate or properly weigh up the benefits and costs of pursuing particular areas of study, to the legitimisation of misappropriated information or resources.

Identification of these harms is hindered by the fact that nowhere is there a clear list of what they are or how they tend to instantiate. Understanding what

these harms are, and how they appear in the wide variety of scientific contexts is, however, critical to responsible scientific research and policy.

This paper attempts to address this gap, laying out a categorisation of the harms of research that cannot be mitigated by focusing on application, and giving examples of them *in situ* drawn from other work in philosophy and science. My goal is to define a useful list that be drawn on by philosophers, scientists, and policy-makers who are working to minimise the harms of scientific research.

2 Harms and Science

The existence of harms[1], either actual or potential, underpins almost all calls for controlling what science can and cannot do (see eg. Pamuk (2021); Stilgoe, Owen, and Macnaghten (2020); US Department of Health and Human Services and others (2013)). We care about stopping science from initiating pandemics, irrevocably altering the environment, changing the genetics of the whole human race, accidentally wiping out species, increasing the deadliness of war, and doing experiments on non-consenting subjects.

Harms, however, are complicated beasts; difficult to find, and difficult to define. In science harms are also widely varied - the potential damages caused by medicine are very different to the potential damages caused by physics. This means that having a clear idea of what sort of harms one needs to look out for, and what to avoid, will depend strongly on the area of study, or even individual scientific project.

It is clear, then, that articulating these harms and what they look like *in situ* is going to be critical to any conversation about how to minimise them.

Still, we will need a working definition of harm. Here, I'm going to use the

[1]There is some reason to think that we might fruitfully replace the term 'harms' here with 'wrongs', and thereby capture these intuitions in a slightly different way. I have no particular reason to select one or the other for this paper but acknowledge that there may be useful distinction to be drawn on in future work on this topic.

term rather loosely. Harm, for this paper, will include any negative impact the research has on individuals or society. In this, I have in mind a broadly comparative account of harm; "that to suffer harm is to be put into a certain sort of comparatively bad state - a state that is worse for one than some relevant alternative state." (Hanser, 2008, p.1) That said, I don't think any of the categories of research harm we discuss below will hinge too strongly on taking this particular approach. Here, harms will include those states which are physically instantiated as damages as well as those which simply increase the risk of damages. Thus, one might be harmed by physically being in a car accident, but also by having someone tamper with your breaks such that your chance of being in an accident increases. You might also be harmed if social or legal structures are altered in ways that will give you worse outcomes if you need to rely on them (the law now require the death penalty for anyone found at fault for a car accident, or anyone who deliberately hits you with their car now has no legal or social obligation to check if you're okay or help pay for injuries you sustain). Obviously the inclusion of some of these under the category of 'harm' is controversial (Placani, 2017), but for the sake of this paper we need a practical and pragmatic understanding of what types of negative impacts research might have on individuals and society. Thus, while there is a rich debate to be had about the status of things like 'potential harms' it seems obvious to me that if science creates a more racist society, even if that racism hasn't yet had a chance to directly harm an individual, a harm has occurred.

3 A Failure to Recognise Research Harms

We recognise that science causes harms, and have a number of tools in our collective toolboxes for preventing, managing, or minimising these harms. I want to suggest, however, that there are a number of ways in which our approaches

to recognising these harms miss the mark.

First, there is a tendency within science and policy of thinking that the harms of science are best dealt with at the application stage, as research moves out of the lab and into the real world. We see this throughout the field, ranging from the wording of major moratoriums, which often ban or restrict the use of knowledge, while explicitly allowing research to continue [2], to the prevalence of dual-use[3] research as a central concept in biomedical and biology policy. It is even visible in philosophy, where there is an extensive literature on why we ought to restrict application rather than research

Consider, for example, Agazzi (1989), as cited in Johnson's *Forbidden Knowledge and Science as Professional Activity* (1996),

> "To put it briefly: while it is in principle morally acceptable to know everything, and there are no *morally prohibited truths*, not everything that can be *done* is acceptable, and there are *morally prohibited actions*" (Agazzi, 1989, p.206)

Pro-Knowledge arguments often then make the slide between knowledge and research, arguing that because knowledge is the product of research, and that research is distinct from application, we ought not restrict research. Johnson (1996) recognises this slide, for example,

[2]Examples of this include human genetic modification, "To be clear, our proposed moratorium does not apply to germline editing for research uses, provided that these studies do not involve the transfer of an embryo to a person's uterus." (Lander et al., 2019), geoengineering, "An International Non-Use Agreement on Solar Geoengineering would not prohibit atmospheric or climate research as such, and it would not place broad limitations on academic freedom. The agreement would instead focus solely on a specific set of measures targeted purely at restricting the development of solar geoengineering technologies under the jurisdiction of the parties to the agreement." (Biermann et al., 2022), and autonomous military drones, "In fact, the ban would apply to development only of fully autonomous weapons, that is, machines that could select and fire on targets without meaningful human control. Research and development activities would be banned if they were directed at technology that could be used exclusively for fully autonomous weapons or that was explicitly intended for use in such weapons." (Docherty, 2016)

[3]Dual use research is "research which can be reasonably anticipated to provide knowledge, information, products, or technologies that could be **directly misapplied** to pose a significant threat with broad potential consequences" (National Institute of Health, 2019). (emphasis mine)

> "Implicit in Agazzi's view is a distinction between knowing and doing. Science is equated with knowing and knowing is understood to be passive. Since knowing is passive, it need not and should not be controlled. On this view, only when a step is taken beyond science do we have something powerful, something that can affect the world and be dangerous. The step is from knowing to doing, from knowledge to using knowledge." (Johnson, 1996, p.207)

Alternatively, we have pro-research arguments like that made by Gärdenfors in his 1989 paper *Is There Anything We Should Not Want to Know?*;

> "The upshot is that controlling fundamental research is not only unwanted but well nigh impossible because it presumes that we can foresee major breakthroughs. In the cases where we suspect that research will lead to dangerous technologies, we should try instead to control the technology. This may be extremely difficult, but it is not impossible" (Gärdenfors, 1989, p.8)

Or arguments in favour of academic freedom, as defended by groups like the AAUP;

> "Teachers are entitled to full freedom in research and in the publication of the results, subject to the adequate performance of their other academic duties. . ." (Kreiser, 2001)

Even where there are calls for controls on research, as with Dummett (1980), these are often grounded in application harms;

> "[o]f the scientific research carried out within any given future period, much of it will have applications, some of them quite unexpected, and that, of these applications, most of those that yield unqualified

> benefits for mankind will either be unexplained or, at best, used to enhance the lives only of people in the wealthy nations, while some will, for certain, be used to create as yet unimagined dangers and horrors. (Dummett, 1981, p.291)

In the same paper he claims that "knowledge is an intrinsic good" (p.293), and that research without impact wouldn't be subject to such restriction.

Second, when we do consider the harms done by research, we often have too narrow a concept of what these harms might look like, with the vast majority of the literature focusing on harms done to test-subjects, which comprises only a small subset of the broader harms research can cause. This lack of awareness, in turn, means we are often unaware of how particular research does damage, or are unable to articulate these harms clearly even where we might see them occurring.

I will set aside the first issue of use vs. knowledge here as the subject of future work, in favour of unpacking the second.

Let us start with the obvious: We talk about the harms of science all the time. Philosophers of science have generated a rich set of conversations about how science isn't done in isolation, and that it has social and personal impact on the world around it. When it comes to talking specifically about the harms done at the research stage, however, these conversations often have a very narrow scope. In particular, when we talk about exceptions to academic research freedom the most common case discussed is that of research which causes direct harm to test-subjects,

> "Those for whom the advancement of knowledge is a supreme value might believe that, in basic research as distinct from applied science and technology, no subject should be declared off limits. Yet there are clear inhibitions on some kinds of research involving human beings,

and indeed animals." (Edsall, 1975)

This type of harm has out-sized representation in conversations about when and where we ought to restrict research. Giordanno, for example, in his 2012 book on academic freedom, almost exclusively considers direct harms to human subjects from testing when talking about possible exceptions to unimpeded academic freedom (Giordano, Coggon, & Cappato, 2012). It is also the subject of some of the only international restrictions on research itself, including the the Belmont Report (National Commission for the Proptection of Human Subjects of Biomedicaland Behavioral Research, 1978), and the Nuremberg Code (Taylor, 1950).

The control of these harms is also the primary goal of Institutional Review Boards (IRB) - by far the most common way that research harms are currently identified and mitigated. Note, however, that the remit of the IRB is harms done to those being tested - indeed they are specifically forbidden from considering anything outside of this scope:

> "The IRB should not consider possible long-range effects of applying knowledge gained in the research (e.g., the possible effects of the research on public policy) as among those research risks that fall within the purview of its responsibility." (US Department of Health and Human Services and others, 2009, §46.111)

The reason for this focus is many-fold. One primary one, however, is the aforementioned bias we exhibit towards restricting the application of research, rather than the research itself. This, in turn, means we are predisposed to only consider research harms where they are direct and inescapable, as with those caused by testing on morally relevant subjects.

This isn't to say that we don't also recognise the existence of other harms caused by research. Some philosophers, for example, have borrowed the term

'Moral Hazard' from the field of insurance, to note how some types of research, simply by being done, might prevent us from taking other, better, steps. In the sustainability space, for example, we might worry that research into carbon capture means that the public stops trying to reduce their carbon output because they assume carbon capture will eventually solve the problem.

Similarly, Kitcher's 1997 paper *An Argument About Free Inquiry* demonstrates that certain types of research agendas can, because of the social structures they exist within, cause social harm to particular vulnerable groups by exacerbating biases, even before the research is used (Kitcher, 1997).

These, along with a number of other examples, however, are scattered and highly specific in nature - aimed at talking about the particular harms of a particular case, without clear recognition of the broader patterns of harm they instantiate. Moreover, these patterns of harm are rarely situated in relation to one-another, nor listed in ways that make it easy to identify all of them for a given case.

There are two particularly critical areas where an understanding and clear categorisation of research harms is going to be important. The first, as already discussed, is for mitigation. We cannot take steps to reduce or prevent the harms of research if we don't have a clear understanding of what they are or where they occur. This is made even more difficult if we're required in each case to figure out the harms from the ground up. Much easier, it seems, to be on the look-out for broad patterns that are recurring in research. A categorisation system, then, allows us to give better guides for policy-makers, scientists, and ethicists, as they work to identify these harms.

The second is in the area of scientific hermeneutical injustice, where in the absence of clear categories victims of such harms may lack the ability to clearly conceptualise exactly what types of harms have been done to them, or articulate

those harms to others. This is particularly insidious when victims are told that the research is value neutral, or even necessarily valuable, and that the harms they've experienced must therefore have come from application. This, in turn, may rob victims of the ability to articulate why the harms happened in the first place, and how to prevent it in the future. Even where victims or regulators may be able to articulate the harms that have occurred in this particular context due to research, an inability to place this harm in the broader context of other harms of a similar type may stymie systematic attempts at justice or prevention.

In short, research harms are under-studied, ignored, or even dismissed as non-existent. An inability to see these categories for what they are, or how they instantiate across science significantly impedes our ability to mitigate, control, or even have comprehensive conversations about the harms of science. We need, then, to begin thinking beyond the narrow scope we've been focusing on, and start recognising both the myriad of other harms science does, and the broad categories they fall under. The first step in this will be the construction of clear categories of research harm.

4 An Incomplete Categorisation of Research Harms

The rest of this paper starts to lay down an initial categorisation of the harms that occur at the research state of science. In doing so, I hope to highlight areas of harms that are under-considered, and provide a useful tool for scientists, ethicists, and policy-makers who may be trying to identify research harms, either along broad lines or for specific research projects.

What follows is not, I want to emphasise, a complete categorisation. As with all things science is messy, nuanced, and hard to generalise. I will miss some harms, sometimes very serious ones, but I hope to capture along broad lines a set of the most common harms caused by of scientific research and which, moreover,

cannot be prevented by purely controlling its use.

The harms I will briefly illustrate below are;

1. Harmful Research Practices
2. Inherently Harmful Questions
3. Harmful Diversion of Resources and Attention
4. Moral Hazards
5. Psychologically Harmful Research
6. Socially Harmful Research
7. Significantly Increased Risk
8. Annexation and Perpetuation

This list is incomplete. As with all science, there will also be difficult fringe cases that don't seem to cleanly fit into any one category. What I hope to do here, however, is offer the largest ones which are relevant to the way we talk about, control, and evaluate the impacts of scientific research.

Thus, while there could be a scenario where someone threatens to kill an innocent bystander if a particular line of research is pursued, and this would indeed be a harm of the research, we'll consider that too narrow to include here.

We can broadly categorise the above list into two subgroups - harms that come from the scientific research itself (what I will call acquisition harms), and harms which come from its position within broader society (context harms). (1 and 2) fall into the former category, (3, 4, 5, 6, 7 and 8) into the latter.

The divide I'm drawing here isn't a hard one, harms can be both acquisition and context, and migrate between the two depending on a diverse number of factors. Even so, it is useful to have the language to talk about the difference

between harms that are largely inherent to the act of research or acquisition of knowledge itself, and those which occur because of the social or political context within which it is being done.

It is also worth noting that context harms are, in many ways, related to the reasoning for standard application controls. That is, they arise because the research is coming into contact with the real world. Here, however, I take the harms to be of the kind that warrant consideration specifically because placing bans on application would be ineffective because once the research is done the harms have already occurred, or are almost certain to occur, and as such if we simply try to intervene on its use once the research has been completed we've missed our opportunity.

Finally simply because a particular line of research causes harm doesn't mean we are forbidden from pursuing it. Many things cause harm, and in most cases these harms are outweighed by the benefits. As such, while each of these harms ought to be recognised, I am not advocating for the idea that we have a moral obligation to prevent *all* the possible harms that might arise from scientific work nor halt any research simply because it causes one of these harms. Research, by its very nature as something embedded in society, will have knock-on effects that are both good and bad. If we banned all things which *could* or even *do* cause harm, we would be in a bad state indeed. Scientists and policy-makers do, however, have an obligation to go into their research with their eyes open to possible negative consequences, sometimes because this means that they ought not to do the research, but more often because our ability to minimise damages while maximising benefits necessarily relies on our ability to identify both outcomes, often from as early as possible in the process. In a similar manner, we very clearly already allow research which causes harm or poses a risk to test subjects - experimental medical treatments, for example -

provided it is done in a way that respects someone's right to choose what harms they suffer, and for what goods. We must not mistake a call to be aware of these harms, for a call to never let harm happen. The first is an ethical obligation, the second would devastate science at best, and be impossible at worst.

Let us look, then, at each of these categories in turn, defining them more clearly and seeing how they instantiate in some real-world cases. This last bit is particularly important for two reasons. First, the kinds of harms we are discussing here are relatively loosely defined and, as such, it will be helpful to see examples. Definitions are, unfortunately, often not enough in this kind of scenario to give us a good idea of what these kinds of problems might actually look like *in situ*. Second, science is, as we all know, a relatively messy endeavour - no two cases of the same type of harm are going to look the same. As such, seeing clear examples can help us more easily identify similarities to other cases. As with many things, the application of these categories will involve careful consideration within a specific context. All I can give is a useful guide.

Many of the following categories have been discussed in the philosophy of science literature. I am certainly not the first person to notice that research can cause harm along these lines. What I offer, however, is a way to understand these harms as part of broader patterns that cross-cut science as a whole.

4.1 Acquisition Harms

Acquisition harms are inherent to the research or knowledge itself. That is, they cannot be separated out and will occur under any circumstances where the research is performed.

4.1.1 Harmful Knowledge Acquisition

Research which gathers information via harm.

Research harms of this kind are those where the very act of collecting data does harm. Within this category we can recognise a number of sub-divisions the most salient of them being information derived from the torture or exploitation of human subjects. Knowledge of how long humans can survive in cold water is directly derived from unethical research performed by the Nazis on war prisoners, for example (Taylor, 1950). It also covers cases like the Stanford prison experiment, Milgram experiment and Tuskegee Syphilis study, and is one of the primary arguments used against STEM Cell work, where critics argue that the act of gathering cells causes harm to human agents (albeit controversially). Other regulated research in this category includes experiments that require hurting certain types of animals that humans consider to be of sufficient moral standing, and experiments that involve doing significant damage to the environment (think, for example, of the lingering effects of dropping an experimental atomic bomb at Bikini Atoll).

Beyond these standard examples, however, we also have cases like iron seeding, which involves spreading iron in the ocean to encourage algae blooms that, in turn, capture and sequester carbon from the atmosphere. Research into the viability of this technique for reducing climate change is, however, almost impossible without actually releasing significant amounts of iron into the ocean to see what happens - there are no small-scale experiments which can be done in the lab which would accurately reflect what we would expect to see *in situ* as the algae interact with broader ocean biomes and chemistry. Thus, the act of doing research is almost indistinguishable from deploying it and, given one of the goals of the research is to figure out what sort of disruptions a large-scale algae bloom cause for deep-sea ecosystems, we cannot do the research without, in turn, causing some of the harms we were doing the research to try and avoid.

There are other, lesser, versions of this harm - research that involves tagging

seals, collecting samples from tide pools, and building laboratories, for example. All of these involve some harm (punching holes in seal ears, taking resources from the land, removing wildlife habitats to construct buildings), but we generally judge these harms to be of sufficiently low stakes that the benefits are acceptable. Alternatively, we might think that research such as that done by NASA during the space race, which involved significant risk to the pilots of the experimental craft and cost the lives of a number of them (Niler, 2019), caused acceptable harm because the pilots themselves consented to the program. Still, however, harm was done.

One thing worth noting here is that we should be careful to distinguish research that *causes* harm, from research which *stems from* harm. Not all research on harmful things is, itself, harmful. Differences in socio-economic prosperity between East and West Berlin, for example, certainly involves harmed parties - one side was less well-off than the other - but it is not the research itself which caused these harms. Similarly neuroscientists looking at the brains of individuals with unique or interesting mental damage or impairments are only causing morally significant harm if they themselves have damaged the patients, as opposed to accidents which happened over the course of the subject's normal lives.

4.2 Inherently Harmful Questions

> *Research which asks questions that cause harm simply by being legitimised.*

The second category of acquisition harms considered here are those where simply considering the research question a legitimate one is, by itself, a harm. This kind of harm often occurs where questions are de-humanising the test subject, or imply that there is some legitimacy to questioning certain fundamental rights

or capabilities.

Take, for example, the question “Should women be allowed to vote?". We can fairly easily imagine what a research project which takes this as its central question might look like. It could look at rates of social well-being between countries that do and do not allow women to vote, for example, or survey women to see how informed they are about candidates and policies. The very act of asking the question legitimises it, however. It gives credence to the idea that whether women are deserving of basic self-determination rights is a question we ought to decide based on research evidence, rather than something which is a given right of all human beings, and that it is an open question whether women and men deserve equal moral standing.

Similarly, questions like “Are Asians as intelligent as Europeans" or "Is homosexuality a genetic disease?" imply a legitimacy that dehumanises and can cause serious harm to the groups who are the target of study.

We can draw on work by Basu (in press) , in understanding this particular harm. She characterise such questions in part as being problematic because they treat humans not as people, but as *objects of inquiry*. As she notes, “Perhaps there are some questions that one cannot ask in a way that is compatible with orienting towards others with dignity and treating them as beings that are worthy as respect, i.e., questions concerning whether or not they are deserving of dignity or respect".

As a brief real-world example, in philosophy (a field notoriously all about asking questions) these kinds of harms came to a head several years ago with a protracted and public debate about trans rights and identities. In a joint statement from Minorities and Philosophy UK (2019), the organisation noted that,

> “Not every item of personal and ideological obsession is worthy of

> philosophical debate. In particular, scepticism about the rights of marginalised groups and individuals, where issues of life and death are at stake, are not up for debate."

This isn't to say that conversations about trans identities are forbidden, but they cannot be done in a way that questions the fundamental rights and dignities of those involved.

4.3 Context Harms

Context harms occur because of the context within which the research is done. As such they are contingent on numerous broader factors and whether or not the harms occur will vary with time and place.

4.3.1 Harmful Diversion of Resources and Attention

> *Research that causes you to abandon, deprioritise, or divert resources and attention away from other actions which might mitigate the problem more efficiently or that might help in tandem.*

This harm is defined in part by limitations on funding or on other broader factors like media coverage and the limited focus of the general public. It essentially constitutes the recognition that there are finite resources for science, whether financial or social, and that pursuing or emphasising some research sometimes comes at the cost of not pursuing others.

Consider, for example, the case of AIDS funding.

In 1988 the AIDS crisis was in full swing, and the federal government was preparing to pass the HOPE act aimed at funding HIV/AIDS research, prevention, and testing. At the time there was extensive debate about how much funding ought to be put towards the disease, and it was openly acknowledged that what

funding was given would likely be drawn from other areas of federal disease research grants.

> "Senator Edward Kennedy has publicly expressed concern that the necessary expansion of the AIDS budged might be financed - directly or indirectly - by slower growth or reductions in other health research programs. Although many citizens might prefer that expanded funding of AIDS-related research come from a reduction of nuclear armament or an increase in taxes, the political reality is that trade-offs are more likely to be made with other health research programs." (Hatziandreu, Graham, & Stoto, 1988, p.1)

One of the serious questions facing policy-makers and scientists was how to weigh up the costs and benefits of particular funding models. Inevitably money taken from cancer research would cause deaths, but so too would failure to fund AIDS research adequately.

Along similar lines, there is ongoing debate about the best focus for AIDS research itself. It is entirely possible that building immunity to HIV/AIDS using the modification of human genomes will prove to be overwhelmingly successful, but it is also a long-shot (Lebbink et al., 2017). Conversely, it may be that more traditional vaccines are more promising, and what funding exists ought to be channelled in that direction. There are also questions of what information is needed to even make these decisions, and how to go about getting them. Debate over HIV/AIDS research priorities are a prominent aspect of the field (Bain & Gwain, 2019; World Health Organization and others, 2012) and there are difficult decisions to be made - do you channel your resources into finding a vaccine as fast as possible, or do you work to try and prevent as many infections as possible? What groups do you focus your research on? Which countries are results rolled out in first?

This is perhaps one of the most difficult harms to navigate, for good reason; it requires difficult decisions in the face of uncertainty, and in the case of medical research it is inevitable that some decisions will lead to worse health-outcomes or even death for numerous people.

Of course, this isn't limited to medicine. There are limited resources available to fund research, and governments, universities, and public grant givers are continually making decisions about how to allocate not just these funds, but also how to focus public attention on the most important of them. Do we fund better solar panels, or more efficient nuclear power plants? Should researchers prioritise building better climate models, or figuring out better ways to communicate currently existing ones to the public and policy-makers? Do we fund malaria vaccines, or buying mosquito nets? The former could save millions of lives, but insecticide-treated mosquito nets have shown to significantly reduce the transmission of the disease (Ntonifor & Veyufambom, 2016), such that a reduction in funding would directly kill people, and an increase would directly save lives.

I also don't want to imply that this harm is of a kind which means that we ought to only ever fund research that maximises human health outcomes, or productivity. This kind of reasoning has historically been weaponised against research in the humanities, for example, or more abstract subjects like theoretical mathematics which has no obvious immediate use, to argue that they ought to be de-funded, or that universities are wasting their resources in employing such researchers. I firmly believe such arguments miss the point of academic research more broadly. Still, it remains the case that this kind of harm is a major factor in the decisions of funding bodies, for example, and so I would be remiss to not include it in this list.

4.3.2 Moral Hazards

> *Research which reduces the incentive to prevent or minimise a damage or harm.*

In science this harm often instantiates via things like techno-optimism, which creates a feeling amongst the general public that we will be able to invent, research, or science our way out of current problems, and thus have no need to take any other steps outside of science to solve them. Unfortunately, because research is just that - research, there is no guarantee that these solutions will ever instantiate or even turn out to be viable.

A number of critiques of climate adaption technologies take this route, hinging on the idea that they are a band-aid, rather than a genuine solution, to the underlying problem of human emissions. Solar Radiation Management Geoengineering - changing our atmosphere to reflect more sunlight - is a good example of this debate in the literature.

One major worry for the technology, however, is the Moral Hazard one - that publicising this research, or even simply having it available, will slow down or even prevent action on climate change.

> "[G]eoengineering could be inaccurately perceived as a comprehensive insurance policy against climate change. This misperception could create various incentives that would exacerbate the problems that geoengineering is intended to ameliorate. Individuals might curb voluntary efforts to reduce carbon emissions. Fossil fuel consumption and other GHG-generating behaviors might even increase out of a misguided belief that climate change no longer poses a threat. Societies might divert resources away from mitigation toward geoengineering schemes that ultimately prove futile or unworkable. Finally, political

> and financial support for mitigation and adaptation policies might decline." (Lin, 2013, p.678)

Exacerbating this is the worry that not only might this technology inadvertently undermine attempts to reduce emissions, but it might be actively weaponised by large-scale polluters to shift focus away from their responsibilities.

Companies like Exxon, Shell, and Boeing have all expressed support for Geo-engineering, funding research and advocating for the technology (Muffett & Feit, 2019). There are concerns, then, that these companies could push a narrative where geoengineering techniques, are an easy and comfortable solution to climate change that wouldn't require major changes to the way we consume resources.

More generally, moral hazards turn up wherever we might worry about 'technology-fixes' or 'science-fixes' coming into play. Why recycle when there are groups working to use bacteria to break down plastics so they're biodegradable? Why conserve energy when green-power is becoming more prevalent? Why enact international policy to help food shortages in third-world countries when genetically modified super-crops are being developed to be used in these places? This kind of thinking is obviously harmful where it guides our actions on the broad-scale, and where we then continually put-off making more effective but perhaps less pleasant, or more personally costly, changes in favour of hypothetical, or even real, results from science.

4.3.3 Psychologically Harmful Research

Research which causes mental anguish, depression, or anxiety.

One of the most salient examples of this harm stems from one of the most pressing problems of our lifetimes: climate change. In particular, there is a recent wave of articles and think-pieces on the phenomenon of 'climate despair'

(Burke, 2021; Clifford, 2021; Vance, 2021). Defined, roughly, as "a sense that climate change is an unstoppable force that will render humanity extinct and renders life in the meantime futile.", it is also sometimes called 'eco-nihilism', or 'human futilitarianism' (Pearl, 2019).

The existence of such feelings should be unsurprising;

> "Few Americans are confident that humans will reduce global warming. About half (49 percent) say humans could reduce global warming, but it's unclear at this point whether we will do what is necessary, and about one in five (22 percent) say we won't reduce global warming because people are unwilling to change their behavior. Only 6 percent say humans can and will successfully reduce global warming." (Frost, 2019)

This despair is a direct result of learning new information about the state of our planet, and the problems that human-caused climate change is causing. Indeed, it grows worse the more up-to-date one is about the situation, and the more information one takes in. Individuals have reported suicidal thoughts, a hesitancy to have children out of fear for their future, and a desire to cut themselves off from learning new information. (Pearl, 2019)

Climate science research is, then, having a direct and terrible toll on people's psyches. It is causing harm to individuals, many of whom would arguably be psychologically better-off not knowing information about how the climate is changing. This is, unfortunately, an inescapable aspect of research into a world-threatening phenomenon.

Now, obviously this research is also necessary and overwhelmingly beneficial. It is hard, after all, to solve an ill-defined problem. The solution to the problem of climate despair is not to stop doing climate research, but to figure out better ways to help people work through the weight of such knowledge. In these situations

the benefits outweigh the harms. This doesn't, however change the fact that the harms exist.

4.3.4 Socially Harmful Research

> *Research which creates further social barriers or difficulties for a given social group or individual.*

Socially harmful research is research which does social harm, rather than just being harmful because of social factors, unlike the other types of harm included in this paper.

This is the argument that Kitcher makes when he argues that research into the intellectual differences between races, regardless of outcome, harms the socially disadvantaged group (Kitcher, 1997). He grounds this claim in a number of factors, firstly that where there is a political asymmetry, studies which tell in favour of the minority group are largely ignored, while studies which find against them are taken up and used to motivate further harms to the group. Second, that because of the epistemic asymmetries involved in being a member of a society with biases, members of the social group are more likely to believe research which confirms their bias, and disregard research which contradicts it. This is, of course, to not even speak of the fact that researchers themselves are members of society and so, while all due caution may be taken, it is certainly the case that these biases have influenced the outcome of studies and will do so again.

So, for example, if one were to try and study the inherent differences in mathematical abilities between Asian-Americans and African-Americans, any study which showed no difference is likely to have less uptake than studies which confirm public bias and show that Asian-Americans are inherently better at mathematics than their African-American counterparts. As such, there is no

benefit to either group. Either the status-quo stays the same, or it gets worse.

This harm is, of course, also related to the Inherently Harmful Questions category above, although is distinguished by causing harms relative to particular social backgrounds eg. racism and sexism. Questions about the different athletic abilities between people who are left or right handed is, for example, unlikely to cause significant harm to one or both groups in ways that similar research between races or gender will, which in turn is different to asking whether certain races have the right to vote.

4.3.5 Significantly Increased Risk

> *Research which, once done, is extremely likely to be used in harmful ways.*

This harm comes under a number of headings, but is sometimes referred to as 'slippery slope', 'scientific momentum' or 'path dependency' problems, or alternatively being 'locked in' (Callies, 2019). While there has been some discussion of this problem for geoengineering, here we will look at lethal autonomous weapons systems (LAWS) research, which suggests that placing a moratorium on application of the techniques is useless if militaries will simply ignore moratoriums to gain combat advantage. While the research itself doesn't seem inherently harmful, once it exists there is the worry that it is almost guaranteed that groups will use it unethically.

LAWS (also known as Killer Robots) are defined by the Department of Defence as "A weapon system that, once activated, can select and engage targets without further intervention by a human operator. This includes human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation." (US Department of Defence, 2012)

There is extensive discussion in political, philosophical, and social circles about the use of these weapons due to the potentially devastating ethical implications of their deployment. In general, there is broad consensus that they will violate the principle of *jus in bellum* by removing the direct responsibility of killing from the hands of agents with moral standing. It has also been suggested that extensive use of LAWS will make it less daunting for countries to enter into conflict, as it reduces the risk to their own soldiers, as well as that it will lead to an increase in preventable war-crimes as militaries will be disincentivised to include the ability to refuse unethical orders as part of the LAWS programming. For these reasons there have been numerous calls in recent years for LAWS to be banned, both from development and deployment. Here, for example, is the Human Rights Watch,

> "Human Rights Watch and [Harvard Law School's International Human Rights Clinic] are calling on governments to:
>
> 1. Work toward an international instrument prohibiting the development, production, and use of the fully autonomous weapons.
> 2. Develop national policies on the issue, which encompass national moratoria on the development, production, and use of the fully autonomous weapons." (Watch, 2014)

There is a clear connection, in the LAWS case, between research and application. Militaries, after all, have strong pressure to have the best weapons at their disposal. Many militaries are also very well-funded, and are already using semi-autonomous weapons. As such, when a tactical advantage like LAWS are put on the table, the chance of deployment significantly decreases. As Ariel Conn says, "Once this Pandora's box is opened, it will be hard to close." (Future of Life Institute, 2017)

This risk is well known in the literature.

> "Two paradigmatic examples of this category are stratospheric aerosol injection (SAI), which is a radical geoengineering technology that involves reflecting a fraction of incoming sunlight back into space, and killer robots, which are lethal weapons that select and kill targets without human supervision. In both cases, research beyond a certain stage creates the risk of serious and potentially irreversible harm—because of the complexity of climatic processes in the former and the autonomous intelligence of robots in the latter. Both technologies also involve the risk of deployment by rogue private or state actors to intentionally inflict harm, and researchers in both areas have publicly stated that deployment would be a bad idea." (Pamuk, 2021, p.3)

This is a relatively complex harm of research however. Indeed, arguably it isn't a harm of research at all, but a harm of application - something we're avoiding considering in this paper. That said, it is clear that research cannot be disentangled from application in this case - the very act of doing research significantly increases the risk of future harms to human beings. Moreover, this case is distinguished out from other simple application harms by the fact that any ban on application will be almost impossible to police. Even if major militaries could be convinced to give up the tactical advantage LAWS represent, the research acts as a blueprint for rogue states and actors to create such weapons for their own use. Indeed, it seems difficult to contemplate a scenario where, if the research is done, it is not eventually used, no matter how stringent international controls are.

Regardless of whether application actually occurs, however, the research increases the chance of harm and thus, under our understanding of harms in

this paper, is itself doing harm. Indeed, this is a critical harm for researchers to consider as they advance projects - what would happen if someone applied this research. It is for this reason that we ought to be careful when publishing papers which describe how to increase the virility of diseases, build smaller and more powerful nuclear weapons, or modify human genetics for non-health purposes. All of these fall into the category of research which, once available, is unable to be put back in the box and which, as a direct consequence of availability, has significantly increased the chance of serious future harms.

4.4 Annexation and Perpetuation Harms

> *Research which ignores or delegitimises the rights of individuals or communities to things of value to them.*

Science has a torrid history when it comes to research done on, or around, minoritised or oppressed groups, particularly in areas with a history of colonisation and economic or scientific exploitation. We've already discussed some of the ways that scientific research has done this - the Tuskegee syphilis experiments, for example, or research into relative intelligence between races. One important category not yet captured, however, is that caused to individuals or communities who have their identities, cultures, resources, or other things of value to them appropriated by science for research purposes and, through this, have their own rights to it eroded.

Let me begin with what I take to be a helpful analogy. Imagine someone wandered into your house and told you that your fridge was now common property. They then raid your fridge for food which they take back to their house, make a delicious dinner with it, and feed their friends. Their friends then start to also show up, raid your fridge, and make their own meals. In fact, your fridge starts to be treated by the community as their own, no matter how

much you protest against it. You may even be told that you're in the wrong for wanting to keep people out of your fridge; see what wonderful benefits there are to the community when it's open! How could you be so selfish as to deny access to those who can make such lovely meals? You might, in fact, get to have some of those meals yourself at some point. It's clear, however, that someone has simply appropriated something that you have legitimate ownership of - the food you bought and had plans for in your fridge - and created and perpetuated an expectation that others can use it without your permission.

There are, I think, two types of sub-harms happening here. The first harm, call it *annexing harm*, comes from the initial annexing of your fridge - someone essentially stole your rights to something you own. The second, call this *perpetuation harms* comes from the continued use of the fridge by others - the more they use it and tell their friends about it, the weaker your claim over your fridge becomes as it shifts from 'private' to 'public' in the minds of the community you reside in. The latter is also clearly a harm regardless of whether those who use your fridge are aware that the fridge was initially stolen from you.

Science can, and has, done both of these sub-harms. On the annexing harm side, we need only look to the history of medicine where western researchers have a long history of appropriating traditional medicines, claiming it as their own, and treating all subsequent discoveries as their own property to be commodified and sold. Alternatively, we might look to examples like that of Henrietta Lacks, who's genetic code was taken without her permission (Wolf, 2011), or indigenous communities who have their genetic code taken and used in research without consultation (Boyer et al., 2011; Harry & Kaneche, 2006).

Perpetuation harms come in the form of other researchers continuing to use the HeLa cell line for their work, each subsequent use normalising the use of Lacks' genetic code as an object of research. This, in turn, harms her

both because it continues to deny her family the right to control their own genetic code, and because it actually erodes their claims in the eyes of the public. The initial annexing doesn't, of course, have to have been by or for science. Take, for example, the Thirty-Meter Telescope's construction in Hawai'i. Here, the US government, acting as a colonial power, took the islands from the indigenous owners and proceeded to use the land for their own purposes (Silva, 2004). Eventually the US government gave permission for the construction of telescopes on Mauna Kea - the largest mountain in Hawai'i and an excellent site for astronomy research. Scientists, in building these telescopes, and doing research at the sites, however, gave tacit legitimacy to the idea that the US government had the right to make such decisions, and perpetuated the harms done to native Hawai'ians who had their land stolen. Indeed, the construction of telescopes has become the focal point of protests against the US's treatment of native Hawai'ians and their culture, and very specifically about whether the government has the right to give permission for the building of these structures on a culturally significant site (Kahanamoku et al., 2020). Scientists, in pushing for further construction, are then perpetuating the initial colonial harms done to the community, and eroding their claims to their traditional lands.

These harms are often hard to see, and steeped in the history and culture of the harmed groups, particularly because they are often done against groups who lack representation or a voice because of their position within societies. As such it is one that needs to be carefully navigated since Annexation and Perpetuation harms can be particularly damaging to already vulnerable groups.

5 Conclusion

Scientists have a duty of care towards the individuals and public who are impacted by their work. To this end, the list above is designed to help identify the possible

harms of research and how they might be minimised. It should not, however, be thought of as a check-list. First, because this is far from a complete listing of the potential harms scientific research can cause, and second because science is a diverse and complex field, and so any such list will need careful and nuanced reflection when applied. Instead, we ought to think of papers like this one as an opening dialogue and useful tool to help advance the project of socially responsible science.

The existence of the categories developed in this paper can, however, help us articulate exactly what kinds of patterns to look for, identify the root causes of harms, and give us insight into how we ought to minimise them if, indeed, we can. They are important, then, for scientists and policy-makers, not to mention those affected by these harms who can use these categories to clearly convey the kinds of wrongs being done.

In recognising these distinctions, we are also better placed to design and implement mitigation strategies. Acquisition harms, for example, are difficult to eliminate, and as such if we judge an acquisition harm significant enough to warrant a ban it is likely going to be a permanent one. In contrast, because context harms are reliant upon broader social norms or structures, it is likely that mitigation techniques aimed at them will be temporary, or at least contingent on the continued existence of the norms and structures. Such controls are likely to come with time limits or particular conditions upon which they are lifted eg. that a new law is implemented, or that further studies are done. It is also clear that different types of harms within these categories make a difference to how we mitigate them. Psychologically harmful research is likely to need to be addressed via new styles of communication, or by changing the underlying message, for example. Moral hazards, on the other hand, are going to require changes in the motivations of the actors, or may entail restricting what type of

research is pursued and how that is conveyed to parties with the power to make decisions over how we act. Put briefly; we cannot mitigate or make informed choices about harms if we don't fully understand what they are.

Part of this dialogue also involves considering what research to pursue, and how to go about pursuing it. We certainly shouldn't avoid research simply because it causes harm, anymore than we ban cars because they cause traffic accidents. The goal of this paper isn't scientific paralysis.

It is also worth noting that we can't simply expect scientists to reckon with these kinds of questions by themselves. Asking our scientific researchers to also be experts in the ethical, social, and political implications of their work along with their own studies isn't just impractical, but doomed to relegate responsible science to a tick mark on funding proposals. Instead, we need to think in terms of broad structural changes that enable communication and collaboration between ethicists, policy-makers, researchers, and communities.

ranging from the importance of academic freedom (Giordano & Harris, 2020; Wilholt, 2010), to the uncertainty of where future discoveries will come from (Gärdenfors, 1989; Pamuk, 2021), to the idea that more research helps inform the safe use of new technologies and techniques(Cioffi et al., 2022; Lander et al., 2019; Reardon, 2014), to a belief that knowledge and research are inherently valuable (Lekka-Kowalik, 1999; McGucken, 1978), or alternatively that knowledge itself is value-free (Agazzi, 1989). This distinction between research and technology, knowledge and use, can be found throughout philosophy, science, and policy both implicitly and explicitly.

References

Agazzi, E. (1989). Responsibility: The genuine ground for the regulation of a free science. *Scientists and their Responsibilities*, 203–219.

Bain, L. E., & Gwain, G. C. (2019). Cardiovascular disease and hiv infection in sub-saharan africa: misplaced priorities in the public health and research agendas? *Frontiers in Cardiovascular Medicine*, *6*, 35.

Basu, R. (in press). Risky inquiry: Developing an ethics for philosophical practice. *Hypatia*.

Biermann, F., Oomen, J., Gupta, A., Ali, S. H., Conca, K., Hajer, M. A., ... others (2022). Solar geoengineering: The case for an international non-use agreement. *Wiley Interdisciplinary Reviews: Climate Change*, *13*(3), e754.

Boyer, B. B., Dillard, D., Woodahl, E. L., Whitener, R., Thummel, K. E., & Burke, W. (2011). Ethical issues in developing pharmacogenetic research partnerships with american indigenous communities. *Clinical Pharmacology & Therapeutics*, *89*(3), 343–345.

Burke, K. (2021). An antidote to climate despair. Retrieved 2022-04-22, from `https://www.americanscientist.org/blog/science-culture/an-antidote-to-climate-despair`

Callies, D. E. (2019). The slippery slope argument against geoengineering research. *Journal of Applied Philosophy*, *36*(4), 675–687.

Cioffi, K., Adelmant, V., van Veen, C., Bingham, L., Deluca, E., Dhillon, S. K., & Evans, B. (2022). Response to request for information. *Center for Human Rights and Global Justice*.

Clifford, C. (2021). Climate change is radicalizing young people — here's what that means and how to combat despair. Retrieved 2022-04-22, from `https://www.cnbc.com/2021/08/30/how-to-stop-climate-change-despair-according-to-professor.html`

Docherty, B. (2016). Making the case: The dangers of killer robots and the need for a preemptive ban. Retrieved 2022-04-22, from `https://www.hrw.org/report/2016/12/09/making-case/`

dangers-killer-robots-and-need-preemptive-ban

Douglas, H. (2021). Scientific freedom and social responsibility. In *Science, freedom, democracy* (pp. 68–87). Routledge.

Dummett, M. (1981). Ought research to be unrestricted? *Grazer Philosophische Studien*, *12*, 281.

Edsall, J. T. (1975). Scientific freedom and responsibility: Report of the aaas committee on scientific freedom and responsibility. *Science*, *188*(4189), 687–693.

Frost, M. (2019). After climate despair. Retrieved 2022-04-22, from https://www.thenewatlantis.com/publications/after-climate-despair

Future of Life Institute. (2017). An open letter to the united nations convention on certain conventional weapons.

Gärdenfors, P. (1989). Is there anything we should not want to know? In *Studies in logic and the foundations of mathematics* (Vol. 126, pp. 63–78). Elsevier.

Giordano, S., Coggon, J., & Cappato, M. (2012). *Scientific freedom*. Bloomsbury Academic.

Giordano, S., & Harris, J. (2020). *The freedom of scientific research: Bridging the gap between science and society*. Manchester University Press.

Hanser, M. (2008). The metaphysics of harm. *Philosophy and Phenomenological Research*, *77*(2), 421–450.

Harry, D., & Kaneche, L. M. (2006). Asserting tribal sovereignty over cultural property: moving towards protection of genetic material and indigenous knowledge. *Seattle J. Soc. Just.*, *5*, 27.

Hatziandreu, E., Graham, J. D., & Stoto, M. A. (1988). Aids and biomedical research funding: comparative analysis. *Reviews of infectious diseases*, *10*(1), 159–167.

Johnson, D. G. (1996). Forbidden knowledge and science as professional activity. *The Monist*, *79*(2), 197–217.

Kahanamoku, S., Alegado, R., Kagawa-Viviani, A., Kamelamela, K. L., Kamai, B., Walkowicz, L. M., ... Neilson, H. (2020). A native hawaiian-led summary of the current impact of constructing the thirty meter telescope on maunakea. *arXiv preprint arXiv:2001.00970*.

Kitcher, P. (1997). An argument about free inquiry. *Noûs*, *31*(3), 279–306.

Kreiser, B. R. (2001). *Policy documents & reports*. ERIC.

Lander, E. S., Baylis, F., Zhang, F., Charpentier, E., Berg, P., Bourgain, C., ... others (2019). *Adopt a moratorium on heritable genome editing*. Nature Publishing Group.

Lebbink, R. J., de Jong, D., Wolters, F., Kruse, E. M., van Ham, P. M., Wiertz, E. J., & Nijhuis, M. (2017). A combinational crispr/cas9 gene-editing approach can halt hiv replication and prevent viral escape. *Scientific reports*, *7*(1), 1–10.

Lekka-Kowalik, A. (1999, 08). "on freedom and limits of scientific inquiry", in: Z. zdybicka et al. (eds.), freedom in contemporary culture, lublin: Rw kul, 1999, 547-561.

Lin, A. C. (2013). Does geoengineering present a moral hazard. *Ecology LQ*, *40*, 673.

McGucken, W. (1978). On freedom and planning in science: The society for freedom in science, 1940-46. *Minerva*, 42–72.

Minorities and Philosophy UK. (2019). Joint statement in response to the aristotelian society talk on 3rd june 2019. Retrieved 2022-04-22, from `https://www.mapforthegap.org.uk/blog/joint-statement-aristotelian-june`

Muffett, C., & Feit, S. (2019). Fuel to the fire; how geoengineering threat-

ens to entrench fossil fuels and accelerate the climate crisis. *Center for International Environmental Law (CIEL)*.

National Institute of Health. (2019, May). Dual use research of concern. *National Institutes of Health*. Retrieved from `https://osp.od.nih.gov/biotechnology/dual-use-research-of-concern/`

National Commission for the Proptection of Human Subjects of Biomedicaland Behavioral Research, M., Bethesda. (1978). *The belmont report: Ethical principles and guidelines for the protection of human subjects of research*. Superintendent of Documents.

Nelkin, D. (1979). The social responsibility of scientists. *Annals of the New York Academy of Sciences*, *334*(1), 176–182.

Niler, E. (2019). How landing the first man on the moon cost dozens of lives. Retrieved 2022-04-22, from `https://www.history.com/news/first-man-moon-cost-lives-kennedy-armstrong`

Ntonifor, N. H., & Veyufambom, S. (2016). Assessing the effective use of mosquito nets in the prevention of malaria in some parts of mezam division, northwest region cameroon. *Malaria journal*, *15*(1), 1–8.

Pamuk, Z. (2021). Risk and fear: Restricting science under uncertainty. *Journal of Applied Philosophy*, *38*(3), 444–460.

Pearl, M. (2019). 'climate despair' is making people give up on life. Retrieved 2022-04-22, from `https://www.vice.com/en/article/j5w374/climate-despair-is-making-people-give-up-on-life`

Placani, A. (2017). When the risk of harm harms. *Law and Philosophy*, *36*(1), 77–100.

Reardon, S. (2014). Viral research moratorium called too broad. *Nature News, 23rd October*.

Silva, N. K. (2004). *Aloha betrayed: Native hawaiian resistance to american*

colonialism. Duke University Press.

Stilgoe, J., Owen, R., & Macnaghten, P. (2020). Developing a framework for responsible innovation. In *The ethics of nanotechnology, geoengineering and clean energy* (pp. 347–359). Routledge.

Taylor, T. (1950). *Final report to the secretary of the army on nuernberg war crimes trials under control council law no. 10* (Vol. 10). US Government Printing Office.

US Department of Defence. (2012). Dir 3000.09, autonomy in weapons systems.

US Department of Health and Human Services and others. (2009). Code of federal regulations: Title 45, public welfare; part 46, protection of human subjects. *Washington, DC: US Dept of Health and Human Services*.

US Department of Health and Human Services and others. (2013). *United states government policy for institutional oversight of life sciences dual use research of concern.* US Department of Health and Human Services Washington, DC.

Vance, E. (2021). What to do about climate despair. Retrieved 2022-04-22, from `https://www.nytimes.com/interactive/2021/07/23/well/mind/well-climate-anxiety-action.html`

Watch, H. R. (2014). Advancing the debate on killer robots: 12 key arguments for the preemptive ban on fully autonomous weapons. *Available from: https://www.hrw.org/sites/default/files/supporting*$_r$*esources/advancing*$_t$*he*$_d$*ebate*$_8$*may2014*$_f$*inal.pdf*.

Wilholt, T. (2010). Scientific freedom: its grounds and their limitations. *Studies in History and Philosophy of Science Part A*, *41*(2), 174–181.

Wolf, J. H. (2011). *The immortal life of henrietta lacks.* Oxford University Press.

World Health Organization and others. (2012). Who priority research questions for tb/hiv in hiv-prevalent and resource-limited settings. 2010. *Available*

from: whqlibdoc. who. int/publications/2010/9789241 500302_ eng. pdf .