

Penultimate draft. Please cite final version forthcoming in *Res Publica*

The Duty to Edit the Human Germline

Parker Crutchfield

Western Michigan University Homer Stryker M.D. School of Medicine

Abstract

Many people find the manipulation of the human germline—editing the DNA of sperm or egg cells such that these genetic changes are passed to the resulting offspring—to be morally impermissible. In this paper, I argue for the claim that editing the human germline is morally permissible. My argument starts with the claim that outcome uncertainty regarding the effects of germline editing shows that the duty to not harm cannot ground the prohibition of germline editing. Instead, if germline editing is wrong, it is wrong because it violates a duty to protect. However, we also have an epistemic duty to gather evidence regarding the effects of editing the human germline which overrides any moral duty to protect future generations. Thus, we have a duty to gather evidence regarding the effects of editing the human germline, which is to say that we have a duty to edit the human germline.

Keywords: germline editing; duty to non-maleficence; duty to protect; epistemic duty.

1. Introduction

It is now possible to edit the DNA of the human germline using CRISPR, a scientific technique of manipulating DNA. Many people find the manipulation of the human germline—editing the DNA of sperm or egg cells or embryos such that these genetic changes are passed to the resulting offspring and potentially their descendants—to be morally impermissible. There is widespread public opposition to the practice, especially for some interventions. (Houtman et al. 2022) And while there is some consensus among researchers that there need to be stringent restrictions and proscriptions upon the practice, they also generally recognize that there are limited circumstances in which it may be acceptable. In this paper, I argue that it should be permitted, at least in some circumstances, because we have an epistemic duty to gather evidence regarding its effects. This position, however, does not preclude regulation of the practice. And it is not to say that it is permissible in many circumstances. The range of circumstances in which it is morally permissible may be quite narrow, limited to, for example, well controlled and transparent research until the duty to edit the human germline is discharged, at which point the practice ought to be stopped or expanded.

The argument proceeds in two stages. I establish in the first stage that the prohibition on human germline editing cannot be grounded in the idea that such practices violate a duty to non-maleficence. Instead, if there is a morally justifiable prohibition, it is because such conduct violates a duty to protect. Even if editing the human germline violates this duty, it doesn't follow that germline editing ought to be prohibited, because other duties may override the duty to protect. In the second stage, I identify one such duty. I argue that we have a *pro tanto* epistemic duty to gather evidence regarding editing the human germline. This epistemic duty to gather evidence conflicts with the duty to protect. This conflict between the duty to gather evidence and the duty to protect is not uncommon. Sometimes the duty to protect overrides the duty to gather evidence and sometimes the duty to gather evidence overrides the duty to protect. I identify the circumstances in which the duty to protect overrides the duty to gather evidence and argue that in the case of editing the human germline, the duty to protect fails to exhibit these circumstances. Thus, the duty to protect fails to override the duty to gather evidence regarding the effects of editing the human germline. The weight of the epistemic duty recommends editing the human germline.

In the course of arguing for these conclusions, I mean “edit the human germline” to mean editing an embryo such that the embryo is intended to fully gestate and become a human being and who may reproduce themselves. Already, human embryos undergo editing; but they are destroyed at or before 14 days of gestation, the point at which it is no longer possible for twins to develop. Since this is already legally permissible, what I am arguing, on the basis of its moral permissibility, is that this cutoff point for legal permissibility should be removed and that the embryo should be allowed to gestate and eventually result in the birth of an infant (others have already argued that it should be extended).(Appleby and Bredenoord 2018; Nuffield Council on Bioethics 2017; Hurlbut et al. 2017) I am, however, indifferent, at least for the present purpose, to the moral status of the embryo or eventual fetus. Let us suppose that the moral status of a 10 day old embryo is the same as a 14 day old embryo, which is the same as 39 week old fetus.

The paper proceeds as follows. I first offer a series of arguments that a duty to not harm others—the duty to non-maleficence—cannot ground the proscription of editing the human germline. I then argue that a better fit for such a proscription is the duty to protect. I allow that editing the human germline violates this duty to protect, but this allowance ends up being irrelevant. In the subsequent sections I introduce the epistemic duty to gather evidence regarding the effects of the human germline and claim that satisfying it trumps the duty to protect. Even if editing the human germline violates the duty to protect, doing so satisfies the more important epistemic duty. The result is the conclusion that editing the human germline should be allowed; research policy should be set such that the duty to edit the human germline can be satisfied. Of course, I am not the first to draw such a conclusion.(Sparrow 2021; Douglas and Devolder 2021; Gyngell, Douglas, and Savulescu 2017) But others have focused primarily on the moral arguments in favor of the practice, while the present argument is epistemological and thus, a novel path to the moral permissibility of editing the human germline, clearing the path for its legal permissibility.

2. The Duty to Non-Maleficence

Every person has a duty to not harm other people. This duty is a *pro tanto* duty. That a given action harms another person is always a reason to not perform that action, but sometimes other factors override the duty to not harm. Satisfying the duty to non-maleficence appears to ground

the proscription of editing the human germline. Much of the concern about the practice is that it may harm both the person who results from the edited embryo and that person's descendants. In a recent systematic review of reasons authors have offered against editing the human germline, by far the most commonly offered reason was that the possibility of off-target effects was a safety risk for the child and subsequent generations.(Van Dijke et al. 2018) The second most frequently offered reason was that the safety risks are themselves unknown. Or, more generally, the worry is that editing the human germline will cause bad things to happen to future people, either the people whose genes are intervened upon or those people's descendants. If we knew that editing the human germline wouldn't cause harm, so goes the reasoning, then it may be permissible. But we don't know that, so we can't do it.¹

2.1 Uncertainty of Outcome

There are three possible ways one might know, or believe with a high degree of confidence, that editing the human germline harms. One is by direct observation of the reports of people who have resulted from editing the human germline. Only two such people have been recorded, and those two people are still very young, certainly too young to tell people whether they've been harmed (presuming for now that it's conceptually possible to attest to being harmed by germline editing). Obviously, if the practice continues to be prohibited then it will continue to be impossible to know that germline editing harms by asking the people who have resulted from edited embryos. So, this is not a method by which one could acquire justification that human germline editing harms the people who result from edited embryos.

¹ In the background is the familiar non-identity problem (Parfit 1984; Kavka 1982; Boonin 2014), or the problem of how to account for wronging future people. A sufficiently detailed discussion of the non-identity problem and its implications for the present argument is too far afield for the present medium. Among the relevant implications, however, are whether germline editing affects the identity of future persons (Omerbasic 2018; Douglas and Devolder 2022; Sparrow 2021) and whether germline editing might harm the resulting person (Harman 2009; Gardner 2015; Boonin 2014; Rivera-López 2009). Detailed discussion of these issues doesn't cut for or against the present argument, beyond the recognition that in the absence of good reasons to think that germline editing is harmful to future persons, opponents of germline editing must confront the non-identity problem, if their opposition is for harm-based reasons.

A second way to acquire such justification is scientifically. But at this point this is also impossible. Science practices exclusively in the domain of descriptive properties. To claim that science can make discoveries about normative properties such as harm is to conflate facts and values. Though some people argue that moral properties are admissible contents of perceptual experience, directly perceiving moral properties in a controlled, scientific manner is far out from accepted scientific practices. (Werner 2020; 2016) Since the potential harms associated with human germline editing can't be directly observed, if they exist they must be inferred. And since there are no people available to observe to support this inference, the inference must be based on observations of a model. That is, the scientific justification to believe that editing the human germline *will* cause harm must proceed by extrapolation from model to target.

The epistemology of this extrapolation shows that it can't support the conclusion that human germline editing harms people. Suppose that the model used is either a human embryo not intended for full gestation or an animal model, such as a non-human primate that results from an edited embryo. Given the use of these models, it is not possible to observe harm in them and then extrapolate to the target population, namely humans whose embryos have been edited. The use of a human embryo, even if we suppose it has the same moral status as an infant or fully developed adult, lacks the target property of interest, namely, harm. Harm is a matter of well-being, and embryos have none. So, using them as a model can't justify an extrapolation to claims about how editing the human germline harms people. Animal models may exhibit something like the target property, but they also won't work for extrapolation. The very use of an animal model for this research requires the assumption that the animal's moral status is lower than that of a human's (otherwise there would be little reason to use the animal rather than a human). If one then observes an off-target effect in that animal and infers harm to the animal, that inference is not available to make about humans. The off-target effect-causes-harm inference can't be made about humans because what counts as harm for a being with one moral status is very likely different from what counts as harm for a being with greater moral status. The assumption

required to conduct the research—animals have lower moral status—undermines its purpose, which is to draw inferences about harms to humans.²

A third way one might support the idea that editing the human germline will harm the humans who result from it is simply by assumption: we *know* that editing the human germline will cause disability and that diseases harm people, so when gene editing results in disability and disease, we know that it has harmed the person.

However: (i) disease is often defined in terms of harm, so it won't work to then identify harm by pointing to disease; (ii) disability does not necessarily harm a person, as the vast literature on disability ethics indicates (though there are obviously some disabilities that do imply harm); (iii) humans can have very high quality lives even in the presence of significant disability and disease; and (iv) as compared to not existing, living with disability and disease is arguably better.

There is thus no way to say that the decision to edit the human germline *will* harm the resulting person. There is no way to access that outcome with the degree of certainty required to say that it *will* happen rather than merely *may* happen. More epistemic (rather than moral) justification is needed. For now, the basic sciences lack epistemic access to the moral properties associated with editing the human germline.

None of the above is to assert that violations of the duty to non-maleficence can occur only in the context of outcome certainty, however. For example, performing an act that is known to harm a person 5% of the time would still violate the duty to non-maleficence.³ But the outcome ignorance in human germline editing runs much deeper. Not only are we all ignorant about the actual likely outcomes of human germline editing, we are also ignorant of the likelihoods themselves. For some biomedical interventions, one might be able to have a high credence in a particular likelihood—one *knows* what the risk of harm is, even if it known to be low. But in editing the human germline, no one is in a position to have high credence regarding

² Even if the moral property could be observed in the model, the extrapolation to humans would be unwarranted. There are a few accounts of the epistemology of extrapolation, but none of them would warrant the extrapolation, because the difference in moral status, in the case of animal models, is a causally relevant difference. See LaFollette and Shanks (1993, 1995), Steel (2008), and Cartwright (1989) for detailed accounts of the epistemology of extrapolation. On none of these accounts could one justifiably extrapolate harm to humans from observations of harm in models.

³ Thanks to an anonymous reviewer for suggesting this point.

likelihoods of harm. No one can claim with confidence the *likelihood* of any particular harmful outcome. Such confidence can only be established empirically, and the observations are so far insufficient to establish not only the potential harms, but also their likelihoods. This is unlike other instances in which the outcome is uncertain but one can have some confidence in the level of risk of harm. Editing the human germline is not like driving on the wrong side of the road or firing a gun into the air, which may plausibly be proscribed by the duty to non-maleficence. While performing the latter acts one might be uncertain of what the outcome will be, one can be confident in the nature of the possible harms and their likelihoods. This is because these harms have been observed directly and indirectly via extrapolation. But such is not the case with human germline editing, about which we have significant higher-order ignorance. This higher-order ignorance undermines the idea that human germline editing violates the duty to non-maleficence.

4

The duty to non-maleficence can't ground the proscription of editing the human germline, regardless of what account of harm one adopts. Even if one adopts a non-comparative account of harm, such as one of the many threshold accounts according to which harm is a matter of falling below a threshold of well-being, non-maleficence cannot ground the proscription upon editing the human germline. The problem with using non-maleficence as grounds for proscription lies in the epistemology of harm, not the morality of it. Threshold accounts, or other non-comparative accounts of harm, only vary according to the metaphysics of harm, not the epistemology of it.

2.2 *Not Not Wrong*

The above argument *does not* support the claim that editing human germline is not wrong. It only supports that if it is wrong, it can't be wrong in virtue of violating the duty to non-maleficence. Nor does the above argument support the claim that editing the human germline won't harm future people. It might. Whether we know it does depends on identifying the specific effects upon the material and psychological interests of human beings. Alternatively, human germline editing might not harm anyone ever. It's this very uncertainty that undermines the

⁴ If this reasoning is accurate, then the higher-order ignorance is an excusing condition.

notion that non-maleficence grounds its wrongness. But there are other ways actions can be wrong, other duties that editing the human germline might violate.

Above, I claim that embryos may have full moral status, but that I am neutral on this point. But suppose that it is right that embryos have the same moral status as you and I. If that's correct, then presumably an embryo has whatever rights we have. And we have the right not to have our DNA edited without our consent. Thus, when an embryo has its DNA edited, its rights are violated, which may constitute a wrong, independent of whether the embryo and resulting person are harmed.

Nothing I am claiming rules out this possibility. Editing the human germline may be wrong for other, non-harm-based reasons. The typical suggestion for how an act might be wrong that is not in virtue of whatever harm that act causes is that the act violates a right. It is compatible with my argument that editing the human germline is wrong because it violates the rights of embryos who have full moral status. But for germline editing to be wrong for this reason, it must also be true that embryos in fact have full moral status and that beings with full moral status have a right to not have their DNA edited without consent. Both claims need a strong defense. In any case, supposing one can offer such a defense of these two claims, the proponent of editing the human germline can retreat to the claim that although editing human embryos may be wrong, editing sperm and egg cells remains permissible. These, without a doubt, lack full moral status, undermining any right they might have to have their DNA edited. In that case, the subsequent argument goes through, but just for editing germ cells rather than embryos.

What is justified is the proposition that editing the human germline *may* harm the resulting person. The move from *will harm* to *may harm* is significant, because the duty to non-maleficence doesn't proscribe actions that *may* harm others. If it did, then we couldn't go about our daily routine without violating that duty—satisfying the duty would hog-tie our behaviors. Any proscription must be grounded in some other duty, namely a duty to not do things that *may* harm another person. This duty, as I discuss in the next section, is a special obligation, a duty to protect.

3. The Duty to Protect

Some people have the duty to protect others. Unlike the duty to non-maleficence, which everyone bears at all times toward all others, the duty to protect is agent-relative, a duty that one bears in virtue of the relations they bear toward others. The most obvious example of the duty to protect is the duty of a parent to protect one's child. Satisfying this duty goes well beyond refraining from harming one's child. One must also shield them from possible harms. This, of course, is not a duty that typically holds outside of specific relationships. But a parent must shield their child from various threats and risks (though certainly not all). To the extent that I am able, I must put away the saws and drills instead of leaving them out on the counter; I must hold my child's hand as they cross the street; I must get them the medicine they need when they get sick; I must stand between them and aggressive dogs; I must do lots of things for them to shield them from dangers that I don't have to do for anyone else. When I am older, however, they will have a duty to protect me, though this duty may be much weaker in comparison. When I am old and frail and vulnerable, my child will have to shield me from some limited set of threats, such as my own inabilities to care for myself.

The duty to protect is not unique to parent-child relationships, however.⁵(Goodin 1985) Lawyers have a duty to protect their clients from legal exposure; physicians and dentists have a duty to protect their patients' privacy; corporate boards have a duty to protect their shareholders' stake in the company; states have some limited duty to protect their citizens from invasion; ship captains have a duty to protect passengers; institutional review boards have a duty to protect potential human subjects; institutional animal care and use committees have a duty to protect research animals. The duty to protect is common throughout society. And it provides a solid foundation for the proscription of human gene editing.

⁵ See Goodin (1985) for a detailed account of all special obligations and, among them, the duty to protect. The details of his account of the duty to protect are not necessary for the present purpose, but he argues that one person's vulnerability to another person's actions or omissions triggers a duty to protect that person. A child's vulnerability to their parents triggers for them a duty to protect the child. The strength of this duty varies according to the degree to which one person is vulnerable to another. An elderly parent is vulnerable to their adult child's actions and omissions, which triggers in the adult child a duty to protect those interests of the elderly parent's which are vulnerable to the adult child's actions or omissions. And so on for the other examples. Passengers are vulnerable to the captain; clients are vulnerable to their attorneys; patients are vulnerable to their doctors; children drowning in shallow ponds are vulnerable to bystanders.

Penultimate draft. Please cite final version forthcoming in *Res Publica*

The duty to protect provides a solid foundation for several reasons. The first is that it can be violated even if there is no bad outcome. Suppose I fail to hold my child's hand as we cross a busy street. I have violated my duty to protect even if several cars swerve to miss us and we end up making it safely across. No harm, but I still wronged him. Or suppose a ship captain fails to check the adequacy of life boats before disembarking with passengers. The captain has done something wrong, even if the ship makes it safely to its destination.

In the case of editing the human germline, the duty to protect is a helpful foundation for its proscription, because even if the resulting people aren't harmed, the person who edits the human germline can still be considered to have done something wrong. This set of circumstances seems to account for the case of He Jiankui and the international uproar he caused when he edited the embryos and then allowed them to gestate to twin girls. Though we don't have good information on the well-being of these children, it is rather irrelevant to the wrongness of Jiankui's actions. He did wrong, regardless of whether they were harmed. And the duty to protect can help to explain why. He had a duty to protect the embryos and the future people that result from them, yet he didn't.

The second feature that positions the duty to protect to provide a solid foundation for the proscription of editing the human germline is that for it to be violated, there need not be a specific, identifiable person to whom protection is owed. Consider, for example, the ship captain. She must ensure the safety and seaworthiness of her ship well before anyone boards, or even before anyone even purchases a ticket. The satisfaction of her duty to protect must occur before there is any specific, identifiable person who is entitled to that protection. This feature is useful for the proscription of editing the human germline, because the wrongness of identity-affecting decisions can be attributed to the act (or omission) itself rather than any attribute of a future person.

The third feature of the duty to protect that positions it to provide a solid foundation of the proscription of editing the human germline is that it is a duty that is already incorporated into standard research regulatory compliance. The whole point of an IRB is to protect potential human subjects from risks, as compared to the benefits to society or those that the subject may directly derive. Of course, just as a parent can't protect their child from all risks, an IRB can't protect human subjects from all risks. But they do require that risks be minimized, even for those studies that are greater than minimal risk. This requirement that risk be minimized is the shield

Penultimate draft. Please cite final version forthcoming in *Res Publica*

that protects human subjects from the risks inherent to the participation research. So, the duty to protect is already integrated into the normal process of regulating research.

Does the duty to protect proscribe editing the human germline? Perhaps it does. Plausibly, He Jiankui had a duty to protect the embryos and the future people that result. If he did, then he plausibly violated it, as he failed to protect those embryos and people from significant risks, those risks he was able to shield them from. Just as a ship captain violates the duty to protect when they fail to check the life boats, Jiankui violated the duty to protect the embryos or the people that resulted, and that's why what he did was wrong. Of course, he also had a duty to not harm them. But in editing the embryos, he didn't violate that duty. Similarly, if a researcher edits a 13 day old embryo, then allows that embryo to gestate to a fully developed fetus which becomes a human child, then it is plausible that the researcher has violated their duty to protect. But, again, when they edited the embryo, they did not violate the duty to non-maleficence. Thus, if editing the human germline is proscribed, it is better to ground that proscription in violation of the duty to protect rather than the duty to non-maleficence.

4. Duty to Gather Evidence

W.K. Clifford tells the story of a shipowner who, failing to check the seaworthiness of his ship, sends it on its way across the sea.(Clifford 1877) The ship is in fact not seaworthy and sinks with all hands lost. The shipowner believed the ship was seaworthy, but his failing to support this belief with reasons was not only an epistemic failure but also a moral one. Clifford says that, "It's wrong always, everywhere, and for anyone to believe anything upon insufficient evidence." Clifford's assertion is that the shipowner ought to have had evidence supporting the seaworthiness of the ship. The shipowner's epistemic failure—he ought to have had evidence—resulted in a moral failure—he ought to have diligently protected the lives of those aboard his ship.

Clifford thinks that all of a person's epistemic failures are moral failures. If I believe I have a penny in my pocket, but I lack sufficient evidence for that belief, I have not only failed to meet my epistemic obligations but I have also failed to meet my moral obligations. Even though the belief itself might be inconsequential, this failure erodes my trustworthiness. Later, perhaps

when I testify to something, I hold myself out as a trustworthy source but in fact I am a bad believer.

Not many believe that Clifford is right that all epistemic failures are moral failures. But it is commonplace to think that there are things one ought to do epistemically. And it is true that the moral properties of an act can depend on the epistemic properties of the actor. This is especially true of biomedical research. Introducing a drug or device to the market based on insufficient evidence is the equivalent of Clifford's shipowner.

It is a natural intuition that when one believes something without supporting evidence they have failed to meet a normative standard. But what one ought to have done—their epistemic duty—is less clear. One proposal is that one has a duty to believe all and only true propositions. This is far too demanding, however, and it is counterintuitive. It's not humanly possible to believe all true propositions. And even good evidence can lead one astray, so believing false propositions doesn't entail that one has failed to meet one's epistemic duty. Another proposal is that one's epistemic duty is to, for all propositions one considers, try to believe a proposition if, and only if, the proposition is true. This is better, and so is a third proposal: that one has a duty to fit one's beliefs to one's evidence.(Chisholm 1966; Feldman 1988; 2002)

However, both of these proposals succumb to the problem of doxastic voluntarism. It is common to think that ought implies can (or, as I argue elsewhere, that some relation other than implication holds between oughts and cans).Click or tap here to enter text. If, for example, you ought to believe that the birthday party features a clown, it must be true that whether you hold that belief is under your control. But, others claim, that belief is not under your control. Since when you are sitting at the picnic table looking right at the clown with his red makeup and sharp teeth, your belief that the party features a clown is not up to you; you can't believe otherwise. Since you can't believe otherwise, and ought implies can, it's false that you ought to believe that the birthday party features a clown. There is tension between the intuition that we have epistemic duties, that ought implies can, and that we don't control what we believe.(Hall and Johnson 1998; Steup 2008; 2000; Chrisman 2008)

Though there are other attempts at resolving this tension, one way to resolve it is to note that doxastic voluntarism doesn't hold synchronically but that it does hold diachronically. When you are sitting at the party looking right at the clown, at that moment the perceptual evidence weighs upon you and decisively so. You can do nothing other than believe that the clown is right there

right then. Similarly, suppose right at the moment one cannot believe that an all-knowing, all-powerful, and all-good god exists. It is simply out of her power to believe in God right at that moment. Doxastic voluntarism does not hold synchronically. But over time she does have some control over what she believes. Someone who likes the odds of Pascal's Wager might start going to church every week, build relationships with other believers, and regularly read religious texts. If a person wanted to go from believing that the etiology of climate change is anthropogenic to the belief that the etiology of climate change is due to naturally occurring cycles, he could do things that would cause that belief, such as limiting one's science and news consumption to media that make that claim. This is not to say that these beliefs would be justified. But it is to say that over time we have some control over what we believe. And if that's true, then it can also be true that ought implies can and that we have diachronic epistemic duties.

If there are things we epistemically ought to do, and these obligations are diachronic, then our epistemic duties are satisfied by what activities we perform over time, not necessarily what particular state a person is in. Examples of diachronic moral duties are one's duty to maintain a safe environment for their child, or their duty to drive in a safe manner. These are duties that can be violated in a moment, but their satisfaction occurs over time. Examples of diachronic prudential duties are one's duty to maintain a healthy diet, given one's goal of losing weight, or one's duty to practice the piano, given one's goal of putting on flawless performance. What, then, are our epistemic goals?

The most obvious answer to this question is that the primary epistemic goal is to know, or at least to justifiably believe true things. If this is an epistemic goal, there are actions that one must perform to achieve it. If a person wants to know the proofs for Gödel's incompleteness theorems, she must perform them herself. If she wants to know whether she has the ingredients for a loaf of walnut raisin bread, she must look in the cupboards. If she wants to know whether increasing the dose of a drug increases its therapeutic value but not its risk, she must do the research. Given these goals, these are her epistemic duties.

The common feature of all of these actions is that they are all instances of gathering evidence. If we lack knowledge of something, but we want to know it, then we must gather the evidence that will achieve that goal. Relative to a particular epistemic goal, we have a *pro tanto* epistemic duty to gather evidence. (Hall and Johnson 1998) If we generally have the goal of

knowing what the effects of editing the human germline are, we have a *pro tanto* epistemic duty to gather evidence.

This *pro tanto* epistemic duty to gather evidence regarding the effects of editing the human germline often conflicts with the *pro tanto* moral duty to not edit the human germline.⁶ But this conflict is rather common. To resolve this conflict it is thus useful to look to how other conflicts between these two *pro tanto* duties are resolved.

5. Conflicting Duties

Satisfying the duty to gather evidence regarding the effects of editing the human germline requires permitting such practices, ideally in the context of significant research oversight. If the duty to protect also proscribes the practice, then the duty to gather evidence conflicts with the duty to protect. The conflict between the duty to gather evidence and the duty to protect is not an uncommon one.

This conflict is obvious in the United States Constitution. The Fourth Amendment states:⁷

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

When the state has the epistemic duty to gather evidence, for example when it wants to know whether a crime was committed or whether a particular person has acted illegally, the Fourth Amendment establishes that the duty to protect the individual overrides the epistemic duty. The epistemic duty is *pro tanto*. And when other things are considered, such as the duty to protect the individual from intrusion by the state, the epistemic duty cannot rise to an actual duty, unless other conditions obtain, such as conditions on what evidence has already been gathered. The

⁶ Though not always. If one of the duties has been discharged, then they won't conflict. Or, it could be that no one has the goal of knowing what the effects of a particular intervention are, in which case no duty to gather evidence will be generated.

⁷ U.S. Const. amend IV.

state merely having the goal of wanting to know something isn't sufficient for violating the state's conflicting duty to protect.

A similar conflict is present in the Fifth Amendment of the U.S. Constitution.⁸ It states: No person shall be held to answer for a capital, or otherwise infamous crime, unless on a presentment or indictment of a grand jury, except in cases arising in the land or naval forces, or in the militia, when in actual service in time of war or public danger; nor shall any person be subject for the same offense to be twice put in jeopardy of life or limb; nor shall be compelled in any criminal case to be a witness against himself, nor be deprived of life, liberty, or property, without due process of law; nor shall private property be taken for public use, without just compensation.

The relevant part of the Fifth Amendment is protection from self-incrimination, as no person "shall be compelled in any criminal case to be a witness against himself." Again, the state wanting to know something is insufficient for overriding the duty to protect. In the case of the Fifth Amendment, protection from self-incrimination is nearly absolute, unlike the Fourth Amendment in which the duty to protect can be overridden by the epistemic duty. In the Fifth Amendment, there are no considerations that would move the *pro tanto* epistemic duty to gather evidence to an actual duty. The duty to protect an individual from being forced to incriminate oneself will always override it.

The conflict between the epistemic duty to gather evidence and the duty to protect is also common in biomedical research. When researchers are satisfying the epistemic duty to gather evidence regarding the safe dosage of a drug and either the epistemic duty to gather evidence has been discharged (by having collected sufficient evidence) or the ratio of risks to benefits indicates that the risks outweigh the benefits, the duty to protect overrides the duty to gather evidence. For example, suppose that researchers discover during the course of a clinical trial of a drug that it lacks the benefit that it was thought to have and that it carries significant risks. In this case, the moral duty to protect subjects overrides the duty to gather evidence and the research must be stopped.

When the epistemic duty to gather evidence and the moral duty to protect conflict, it is not always the case that the duty to protect overrides the duty to gather evidence. Consider the

⁸ U.S. Const. amend V.

requirement of a physician to report individual cases of some infectious diseases, including sexually transmitted infections, to state public health officials. The duty to protect the patient grounds confidentiality—the duty to protect the patient implies that doctors cannot disclose a patient’s medical record. But the state also has a duty to gather evidence. The duty to gather evidence in this case overrides the duty to protect the patient.

A similar example is the conflict between a lawyer’s duty to protect her client and the state’s duty to gather evidence regarding crime or fraud. A lawyer is duty-bound to protect her client’s confidentiality except when the lawyer has information that may prevent or rectify the consequences of crime or fraud. In some cases, the satisfaction of the state’s duty to gather evidence obligates the lawyer to disclose the information, or the state’s duty to gather evidence overrides the lawyer’s duty to protect her client. In these cases, the duty to gather evidence overrides the duty to protect.

One might say of the STI-disclosing physician that the epistemic duty to gather evidence is derived from the state’s moral duty to protect the public, and it’s this more fundamental moral duty, not the epistemic duty, that overrides the physician’s moral duty to protect the patient. This is a fair point, but that doesn’t mean that the epistemic duty doesn’t also override the physician’s duty to protect the patient. The state’s epistemic duty can still override the physician’s moral duty, even if the epistemic duty is itself derived from a more fundamental moral duty (though the epistemic duty may still be distinct and independent of the moral duty). If Clifford is right, then epistemic duties are always moral duties. Further, Mark Nelson argues that there are no purely positive epistemic duties, as positive duties are always wrapped in a person’s moral or prudential interests.(Nelson 2010) But one might say the same thing about moral or prudential duties. If you have a moral duty to save a child from drowning in a shallow pond, the satisfaction of this duty requires the satisfaction of epistemic duties. I have argued elsewhere that epistemic duties are logically prior to moral (and all other) duties. For instance, you need to know that various actions will satisfy the moral duty, and to know this you have to have some justified beliefs about how this is likely to occur. That is, you need to first satisfy your epistemic duty before you can satisfy your moral duty. Moral duties require actions, and actions require the satisfaction of epistemic oughts to carry them out.

6. Overriding Duties

The duty to protect sometimes overrides the duty to gather evidence, and sometimes the duty to gather evidence overrides the duty to protect. In the case of conducting research to study the effects of editing the human germline, the apparent epistemic duty to do so conflicts with the moral duty to protect. When duties conflict there are a couple of ways one could determine which one to choose. One way is to have some general umpire principle, such as a principle that says choose the one that will generate the greatest utility.(Sidgwick 2011) In the case of research on editing the human germline, the potential for utility is apparently great. Diseases which cause significant suffering throughout humans' lives may be prevented or cured. Further, the research may lead to not only the prevention or cure of disease and the subsequent suffering, but also to the promotion of health and well-being through enhancement.

The utility of researching the effects of editing the human germline are potentially great, so great that one could make the case that we have a moral duty to pursue it. But for those who think utility is morally irrelevant, an umpire principle that adjudicates conflicting duties according to utility is not helpful. There are other ways to choose between duties. As the duty to gather evidence and the duty to protect commonly conflict, one way to choose between them is to examine other instances of this conflict and compare them to the present conflict.

I claim above that my argument is indifferent to whether editing the human germline violates the duty to protect. Perhaps it doesn't violate it, in which case there is no conflict with the duty to gather evidence and the practice should be permitted. But let us suppose that editing the human germline does violate the duty to protect. If true, there is a genuine conflict with the duty to gather evidence.

Consider first an instance in which the duty to protect overrides the duty to gather evidence. The Fifth Amendment is a clear case of this: the state has an epistemic duty to gather evidence, for example because it wants to know who to hold accountable for fraud, but the duty to protect the right of a person to not self-incriminate is absolute. There is no criminal circumstance in which a person cannot exercise this right.

The Fourth Amendment is less absolute. The state has a duty to protect a person's property from search and seizure, but it has a conflicting duty to gather evidence, such as when it wants to know who to hold accountable for fraud. The duty to protect from search and seizure overrides the duty to gather evidence, but not always. The duty to protect is weightier when the search and

seizure of property is warrantless. The conditions for when search and seizure are warranted are the conditions for when the epistemic duty to gather evidence overrides the moral duty to protect property from search and seizure.

A third instance of the duty to protect overriding the duty to gather evidence is when, in the course of conducting research, through standard data monitoring practices or intermediate observation, researchers discover that the intervention under investigation reliably harms the subjects while conferring no significant benefit, such as in a study looking at different dosages of a drug already known to be beneficial. One of the more famous instances of this conflict is that of Zimbardo's prison experiment, in which, despite gathering evidence relating to human behavior, researchers halted the experiment when the gathering of this evidence required that some of the subjects be harmed.

These instances of the duty to protect overriding the duty to gather evidence have important features in common. First, if the duty to protect is not satisfied, either by being violated or by being overridden, in all three cases harm is likely. Second, the subject of the harm is easily identified. Third, the expected harm that the identified subject will experience is preventable. Fourth, in all of the cases the subjects of the harm are engaged in a trusting relationship with the other parties. The failure to satisfy the duty to protect would violate that trust of those harmed, potentially eroding the concept of trust.

Consider next an instance in which the duty to gather evidence conflicts with, and overrides, the duty to protect. A physician has a duty to protect her patient that can only be satisfied by the physician maintaining the patient's confidentiality. Suppose a person goes to their family physician and is diagnosed with a sexually transmitted infection such as gonorrhea or chlamydia. The duty to protect the patient by maintaining confidentiality conflicts with the state's duty to gather evidence of sexually transmitted infections. The duty to gather evidence in this case overrides the duty to protect, and the physician must break confidentiality and report the condition and identity.

Examples of the duty to gather evidence overriding the duty to protect are not difficult to identify. The state's duty to gather evidence regarding a crime overrides the state's duty to protect privacy interests. Indeed, in most contexts in which an agent has a duty to protect

privacy, there will be a potential conflicting duty to gather evidence the satisfaction of which would require intrusion on this privacy.

There are differences between this case and the examples of the duty to protect overriding the duty to gather evidence. First, there is no obvious harm to the patient whose medical information is disclosed to state public health officials. That is, it is false that there is an identifiable harm that an identifiable person will experience. There *is* an easily identifiable subject of potential harm, but it's an open question whether they will be harmed by the violation of the duty to protect. In the case of the STI-disclosing physician, it's not that the effect of disclosure is uncertain—we can be confident that the patient is unlikely to experience harm. Such is not the case when the duty to protect overrides the duty to gather evidence.

Second, since there is no expected or likely harm, there is no identifiable preventable harm. Third, though a physician disclosing medical information to state public health officials may violate a patient's trust in that physician, part of a person's trust in the doctor-patient relationship is built upon the physician's commitment to the health of others. Thus, the violation of trust may not be significant. We could say the same of the person whose privacy is violated by, for example, a search warrant. Part of the trust that obtains between protector and protectee is that the warrant is issued after some due process. It's not clear that one's trust in the state is diminished by a warranted invasion of privacy.

Whether the duty to protect overrides the duty to gather evidence can depend on whether the effect of violating harms a person and whether trust is undermined. But it also can depend on the value of the evidence the gathering of which would satisfy the duty to gather evidence. In the case of disclosing health information to state public health officials, the value of the evidence is very high. Depending on the disease, gathering evidence about who is infected with what diseases will position these officials to prevent the spread of infection to others, which is of high moral value. If state public health officials weren't going to do anything with the information other than store it in a database, then the value would be much lower, and the duty to protect would override the duty to gather evidence.

Similarly, research is stopped when the value of the evidence the gathering of which would satisfy the duty is low, such as when the evidence already gathered confirms or disconfirms a hypothesis. And although the Fourth Amendment protects people from warrantless search and seizure, the balance between the value of evidence and harm to the person tips in favor of the

duty to gather evidence when the state has “probable cause” to search and seize. Having probable cause makes it more likely that the evidence gathered is going to be of high value, as the expected utility of that evidence is greater (because the probability of it being valuable goes up).

7. A Duty to Gather Evidence

When the duty to protect conflicts with the duty to gather evidence, whether one duty overrides the other is dependent on several factors: the certainty and preventability of the harms resulting from failing to satisfy the duty to protect, the degree to which trust is undermined, and the value of the evidence the gathering of which satisfies the duty to gather evidence. In the case of conducting research on editing the human germline, the duty to protect often conflicts with the duty to gather evidence.⁹

On the one hand, the value of the evidence the gathering of which would satisfy the duty to gather evidence is very high. The evidence gathered could result in the mitigation, treatment, or cure of many diseases, potentially preventing human suffering to a degree similar to that of vaccinations. Further, the evidence gathered may also result in enhancements that promote human flourishing.

On the other hand, the identifiability, certainty, and preventability of harms is very low. Suppose the research in question is research on editing human embryos—organisms with, we are supposing, full moral status, who, like everyone else, lack an entitlement to any particular set of genes. This at least allows us to identify the potential victim of failing to satisfy the duty to protect. But, second, it’s still impossible to identify and confidently predict what the effects of editing the genes of an embryo are going to be. To be confident about how editing the human

⁹ They don’t always conflict, however. The absence of this conflict may be due in part to the absence of a goal of knowing what the effects of a particular intervention are. In such a case, no epistemic duty to gather evidence will be generated. Alternatively, there may be no conflict if the evidence gathered is sufficient to discharge the epistemic duty. The absence of the conflict may also be due in part to the absence of a duty to protect. Goodin claims that the duty to protect is triggered by vulnerability between one agent and another. An embryo, or even a future person, may exhibit such vulnerability. It is thus plausible that any instance of editing the human germline occurs in the context of the duty to protect. But the duty may be so easily satisfied—suppose that there is sufficient evidence that an intervention will prevent severe disease—that there is no conflict between the duty to protect and the duty to gather evidence. Or, in other words, when one or both of the duties is satisfied, there is not conflict between them.

germline affects well-being, researchers would need to be able to observe, infer, or assume the presence of harms, and they are not in a position to do any of those.

Third, given our ignorance about the presence or absence of harms, we are also ignorant about the preventability of harm. We won't be in a position to say of a particular harm that it was preventable until we know what the effects of editing the human germline are. So, the certainty of effect is very low, and the value of the evidence is high. What about trust?

Embryos may have full moral status, but embryos are not capable of engaging in trusting relationships with others. Trusting another person is a mental state, and embryos lack minds. Lacking a mind, there is no trust to be undermined by failing to protect it.

If the conflict between the duty to protect and the duty to gather evidence is adjudicated based on the expected harms, the undermining of trust, and the value of the potentially gathered evidence, then in the case of conducting research on the effects of editing the human germline, the duty to gather evidence overrides the duty to protect.

One might additionally think that some potential edits fail to generate a duty to gather evidence. If true, the above analysis would be incomplete, because I would need to offer criteria for when a proposed intervention generates a duty to gather evidence and when it doesn't. For example, suppose a scientist wants to edit an embryo (or sperm or egg) so that the resulting individual suffers so much that their well-being falls well below any plausible threshold of living a life worth living and no other reason. One might think that this proposed intervention fails to generate an epistemic duty to gather evidence. But I think this is wrong. The person who wants to know what happens when one intervenes in this way does have an epistemic duty to gather evidence. The goal of knowing suffices to generate the epistemic duty, even if its satisfaction is morally wrong. It's just that in these cases, other considerations override the epistemic duty, such as the fact that it would violate the duty to protect according to the criteria identified above. The person with this epistemic duty to gather evidence has an overriding moral duty (derived from the duty to protect) to not carry out the proposed intervention.

But ordinarily proposed interventions on the human germline are not like this. In the ordinary cases, the duty to gather evidence overrides the duty to protect. Since in these ordinary cases there are no other duties that override the duty to gather evidence (e.g., the duty to non-maleficence or the duty to beneficence or some other epistemic duty), the duty to gather evidence regarding the effects of editing the human germline is an actual duty. Thus, we have a duty to

gather such evidence. To satisfy this duty, a regulatory scheme that allows research on the effects of editing the human is warranted.

8. Conclusion

One might simply deny that she has the epistemic duty at all, on the grounds that she doesn't have the goal of knowing what the effects of editing the human germline are. One objection to Clifford's view that he considers is that satisfying epistemic duties demands too much time. His reply is that if you don't have time to gather evidence, then you don't have time to believe. A similar reply is warranted here. If a person denies the epistemic duty, she doesn't want to have beliefs about the effects of editing the human germline. She doesn't get to deny that she has the epistemic duty to gather evidence regarding the effects of editing the human germline and still hold beliefs about them.

If I am right that the duty to gather evidence overrides the duty to protect, then the prohibition of such research is not justified. But to say the prohibition is not justified is not to say that the alternative is to leave research on editing the human germline unregulated. The research must be allowed, and this obligation can be satisfied even if there are limits to the research. It may be reasonable for regulations to permit research based on how much value the edit under investigation is anticipated to produce. This permission may limit research on editing the human germline to research that treats or prevents disease or introduces valuable enhancements. It may prohibit, for example, research in how to make humans fluorescent or other research that is not clearly valuable.¹⁰(Gilles, Schinko, and Averof 2015) That research on editing the human germline should be permitted does not imply that the research should be reckless, fail to minimize risks, or otherwise fail to conform with standard scientific practices and human subjects research protections.

It is possible that the duty to gather evidence could be discharged very quickly, especially if the observed effects of a diverse range of genetic interventions are highly undesirable. Such evidence would increase the certainty of effect and decrease the value of further evidence,

¹⁰ It is certainly conceivable, though, that being fluorescent would be valuable to some people.

Penultimate draft. Please cite final version forthcoming in *Res Publica*

causing the duty to protect to override the duty to gather evidence. It is also possible that it could be discharged quickly because the evidence gathered univocally demonstrates desirable effects. Or progress might be slow, taking generations. But once the gathered evidence suffices to hold justified beliefs about the effects of editing the human germline, the duty will be discharged.

Statements and Declarations

The author declares no competing interests or commitments.

References

- Appleby, John B, and Annelien L Bredenoord. 2018. "Should the 14-Day Rule for Embryo Research Become the 28-Day Rule?" *EMBO Molecular Medicine* 10 (9): e9437. <https://doi.org/https://doi.org/10.15252/emmm.201809437>.
- Boonin, David. 2014. *The Non-Identity Problem and the Ethics of Future People*. Oxford University Press.
- Cartwright, Nancy. 1989. *Nature's Capacities and Their Measurement*. Oxford University Press.
- Chisholm, Roderick M. 1966. *Theory of Knowledge*. Vol. 78. Englewood Cliffs, N.J., Prentice-Hall.
- Chrisman, Matthew. 2008. "Ought to Believe." *Journal of Philosophy* 105 (7).
- Clifford, W K. 1877. "The Ethics of Belief." In *The Ethics of Belief and Other Essays*. Amherst, New York: Prometheus Books.
- Dijke, Ivy Van, Lance Bosch, Annelien L. Bredenoord, Martina Cornel, Sjoerd Repping, and Saskia Hendriks. 2018. "The Ethics of Clinical Applications of Germline Genome Modification: A Systematic Review of Reasons." *Human Reproduction* 33 (9): 1777–96. <https://doi.org/10.1093/humrep/dey257>.
- Douglas, Thomas, and Katrien Devolder. 2021. "GENE EDITING, IDENTITY AND BENEFIT." *The Philosophical Quarterly*, June. <https://doi.org/10.1093/pq/pqab029>.
- . 2022. "Gene Editing, Identity and Benefit." *The Philosophical Quarterly* 72 (2): 305–25.
- Feldman, Richard. 1988. "Epistemic Obligations." *Philosophical Perspectives* 2 (n/a).
- . 2002. "Epistemological Duties." In *The Oxford Handbook of Epistemology*, 362–83. Oxford University Press. <https://doi.org/10.1093/0195130057.003.0013>.
- Gardner, Molly. 2015. "A Harm-Based Solution to the Non-Identity Problem." *Ergo, an Open Access Journal of Philosophy* 2 (20190926): 427–44. <https://doi.org/10.3998/ergo.12405314.0002.017>.
- Gilles, Anna F, Johannes B Schinko, and Michalis Averof. 2015. "Efficient CRISPR-Mediated Gene Targeting and Transgene Replacement in the Beetle Tribolium Castaneum." *Development* 142 (16): 2832 LP – 2839. <http://dev.biologists.org/content/142/16/2832.abstract>.
- GOODIN, R E. 1985. "Protecting the Vulnerable: A Reanalysis of Our Social Responsibilities."

- Gyngell, Christopher, Thomas Douglas, and Julian Savulescu. 2017. "The Ethics of Germline Gene Editing." *Journal of Applied Philosophy* 34 (4): 498–513.
- Hall, Richard J., and Charles R. Johnson. 1998. "The Epistemic Duty to Seek More Evidence." *American Philosophical Quarterly* 35 (2): 129–40.
- Harman, Elizabeth. 2009. "Harming as Causing Harm." In *Harming Future Persons*, edited by M A Roberts and D T Wasserman. Springer Verlag.
- Houtman, Diewertje, Boy Vijlbrief, Marike Polak, Jacqueline Pot, Petra Verhoef, Martina Cornel, and Sam Riedijk. 2022. "Changes in Opinions about Human Germline Gene Editing as a Result of the Dutch DNA-Dialogue Project." *European Journal of Human Genetics*. <https://doi.org/10.1038/s41431-022-01114-w>.
- Hurlbut, J Benjamin, Insoo Hyun, Aaron D Levine, Robin Lovell-Badge, Jeantine E Lunshof, Kirstin R W Matthews, Peter Mills, Alison Murdoch, Martin F Pera, and Christopher Thomas Scott. 2017. "Revisiting the Warnock Rule." *Nature Biotechnology* 35 (11): 1029–42.
- Kavka, Gregory S. 1982. "The Paradox of Future Individuals." *Philosophy and Public Affairs* 11 (2).
- LaFollette, H, and N Shanks. 1996. *Brute Science: Dilemmas of Animal Experimentation*. Brute Science: Dilemmas of Animal Experimentation. Routledge. https://books.google.com/books?id=KmEx_bhnsDQC.
- LaFollette, Hugh, and Niall Shanks. 1993. "Animal Models in Biomedical Research: Some Epistemological Worries." *Public Affairs Quarterly* 7 (2): 113–30.
- . 1995. "Two Models of Models in Biomedical Research." *The Philosophical Quarterly (1950-)* 45 (179): 141–60. <https://doi.org/10.2307/2220412>.
- Nelson, Mark T. 2010. "We Have No Positive Epistemic Duties." *Mind* 119 (473).
- Nuffield Council on Bioethics. 2017. "Human Embryo Culture."
- Omerbasic, Alina. 2018. "Genome Editing, Non-Identity and the Notion of Harm." In *Between Moral Hazard and Legal Uncertainty: Ethical, Legal and Societal Challenges of Human Genome Editing*, edited by Matthias Braun, Hannah Schickl, and Peter Dabrock, 67–81. Wiesbaden: Springer Fachmedien Wiesbaden. https://doi.org/10.1007/978-3-658-22660-2_5.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford University Press.
- Rivera-López, Eduardo. 2009. "Individual Procreative Responsibility and the Non-Identity Problem." *Pacific Philosophical Quarterly* 90 (3): 336–63. <https://doi.org/10.1111/j.1468-0114.2009.01344.x>.
- Sidgwick, Henry. 2011. *The Methods of Ethics*. *Cambridge Library Collection - Philosophy*. Cambridge: Cambridge University Press. <https://doi.org/DOI:10.1017/CBO9781139136617>.
- Sparrow, Robert. 2021. "Human Germline Genome Editing: On the Nature of Our Reasons to Genome Edit." *American Journal of Bioethics*. <https://doi.org/10.1080/15265161.2021.1907480>.
- Steel, Daniel. 2008. *Across the Boundaries: Extrapolation in Biology and Social Science*. Oxford University Press.
- Steup, Matthias. 2000. "Doxastic Voluntarism and Epistemic Deontology." *Acta Analytica* 15 (1).
- . 2008. "Doxastic Freedom." *Synthese* 161 (3).

Penultimate draft. Please cite final version forthcoming in *Res Publica*

Werner, Preston J. 2016. "Moral Perception and the Contents of Experience." *Journal of Moral Philosophy* 13 (3): 294–317. <https://doi.org/10.1163/17455243-4681063>.

Werner, Preston J. 2020. "Which Moral Properties Are Eligible for Perceptual Awareness?" *Journal of Moral Philosophy* 17 (3).