

Target Article

Shared Vulnerabilities in Research

Eric Chwang, University of Colorado at Boulder

The U.S. Code of Federal Regulations governing federally funded research on human subjects assumes that harmful research is sometimes morally justifiable because the beneficiaries of that research share a particular vulnerability with its subjects. In this article, I argue against this assumption, which occurs in every subpart of the Code of Federal Regulations that deals with specific vulnerable populations (pregnant women, fetuses, neonates, prisoners, and children). I argue that shared vulnerability is no exception to the general principle that harming one person in order to benefit another is no more justifiable if the two people have traits in common. Further, shared vulnerability is not a reasonable proxy for any morally relevant desideratum of research, in particular the desire to benefit the worst off, the desire to avoid exploitation, and the desire to use vulnerable populations in research only when necessary.

Keywords: biomedical research, human subjects research, philosophy, research ethics

The U.S. Code of Federal Regulations governing federally funded research on human subjects, Title 45 (Public Welfare), part 46 (protection of human subjects), hereafter referred to as CFR 45.46, assumes that harmful research is sometimes morally justifiable because the beneficiaries of that research share a particular vulnerability with its subjects. In this article, I argue against this assumption. Ana Iltis has also recently argued for a similar conclusion (Iltis 2011), but her argument focuses on children whereas mine applies to all vulnerable populations; she assumes that we should be more protective of sick children than of healthy children, whereas I do not; and she does not discuss various possible objections, which I do. The assumption I scrutinize occurs in every subpart of CFR 45.46 that deals with specific vulnerable populations. The first such vulnerable population is the loosely related group of pregnant women, fetuses, and neonates, dealt with in subpart B; the second is prisoners, as discussed in subpart C; and the third is children, discussed in subpart D. CFR 45.46's guidelines for research on all of these vulnerable populations presuppose, incorrectly I argue, that research on these vulnerable populations is sometimes justifiable partly because its beneficiaries share some vulnerability with its subjects.

It is no coincidence that the objectionable presupposition occurs for all research on vulnerable populations. The Code of Federal Regulations owes this common theme to the 1979 Belmont Report, the ancestor to and theoretical basis of the current CFR 45.46. Part C (Applications), section 3 (Selection of Subjects), of the Belmont Report states, with my emphasis:

When research is proposed that involves risks and does not include a therapeutic component, other less burdened classes

of persons should be called upon first to accept these risks of research, *except where the research is directly related to the specific conditions of the class involved.* (National Commission for the Protection of Human Subjects [Belmont Report] 1979)

No argument is given for this assertion, though presumably it is supposed to follow from some combination of the basic ethical principles that the report articulates: respect for persons, beneficence, and justice. And we can surely understand the Belmont Report's basic motivation: Vulnerable populations should typically be protected from harmful research. Of course, the desire to protect the vulnerable does not mean that we should never recruit vulnerable populations for research—only that research on the vulnerable should be the exception rather than the rule.

This line of thinking leads naturally to the difficult question of specifying those exceptional circumstances in which we may recruit the vulnerable to participate in research, and here I think the Belmont Report, and its descendent, CFR 45.46, get things wrong. Thus, while I agree with the spirit of the Report in thinking that the vulnerable should in general be protected, I disagree with how the exception to the general requirement of protection should be specified, and in this article I argue against the Belmont Report's specification, as operationalized in the Code of Federal Regulations, of a common vulnerability shared by beneficiary and subject.

VULNERABLE POPULATIONS

I begin by describing the areas in CFR 45.46 that make the assumption I challenge.

Address correspondence to Eric Chwang, MD, PhD, Assistant Professor, Philosophy Department, Box 232, University of Colorado at Boulder, Boulder, CO 80309–0232, USA. E-mail: chwang@colorado.edu

Pregnant Women, Fetuses, and Neonates

For pregnant women, fetuses, and neonates, the target assumption occurs in §207 of CFR 45.46. This is the last clause of the section dealing with research on pregnant women, fetuses, and neonates, and it gives guidelines for research that fails to satisfy all other possible justifications for permission, for example, the justification that the research holds out prospect of direct benefit to its subjects. It states that research on these vulnerable populations (pregnant women, fetuses, and neonates) that would not otherwise be permissible might still be permissible, but only if various other criteria are all satisfied. The criterion of interest for the purposes of this article is clause §207b2i: “The research presents a reasonable opportunity to further the understanding, prevention, or alleviation of a serious problem affecting the health or welfare of pregnant women, fetuses, or neonates” (Code of Federal Regulations 2009, §207b2i). (Technically, §207a states the same criterion, with the difference being in who is responsible for its determination, the institutional review board [IRB] versus the Secretary of Health and Human Services as informed by a panel of experts. This subtle distinction and redundancy is irrelevant for the purposes of my argument.)

In saying that §207b2i permits potentially harmful research on a vulnerable group only if others with that same vulnerability are benefited, I have generalized the language of CFR 45.46. Subpart §207 does not discuss benefits to pregnant women, fetuses, or neonates directly; rather, it says that research must “present an opportunity to understand, prevent, or alleviate a serious problem affecting the health or welfare of pregnant women, fetuses, or neonates” (Code of Federal Regulations 2009, §207b2i). I think the generalization from that quoted phrase to benefits is both innocuous and an improvement over CFR 45.46’s position. It is innocuous because the main point of understanding, preventing, or alleviating a serious problem affecting the health or welfare of pregnant women, fetuses, or neonates is clearly to benefit pregnant women, fetuses, and neonates. Yes, there are ancillary benefits, such as the benefits to a pregnant woman’s family, but those ancillary benefits are not the point of the restriction.

My generalization is also an improvement over CFR 45.46’s position. That is because there are benefits beyond those related to health and welfare, and, while CFR 45.46 claims that improvements in health or welfare of members of a vulnerable population can justify harming others in that population, a logically weaker and therefore more plausible position would be that benefits of any kind might do that same justificatory work. Thus, in switching from opportunities related to health and welfare to the more general concept of benefits, I am being charitable to CFR 45.46, in that I attribute to it an even weaker position than it officially holds. Still, I argue that no benefit—whether health-related, welfare-related, or otherwise—can justify research merely because it is reaped by others who are similarly vulnerable to those harmed, and a central set of illustrative examples I use in my argument involve health-related benefits.

Prisoners

The relevant clause in CFR 45.46 regarding prisoners is §306. In effect, all research on prisoners must either have a reasonable chance of improving the health or well-being of its subjects (Code of Federal Regulations 2009, §306a2iv) or else it must study something related to criminality, prisons, or prisoners (Code of Federal Regulations 2009, §306a2i, ii, and iii). Again, this latter restriction is, strictly speaking, not about benefits to other prisoners (guards may also be benefited from studies on prisons, for example), but in practice benefit to other prisoners and other potential criminals is clearly the main point of the restriction.

Children

For children, the Code of Federal Regulations contains two relevant clauses, §406 and §407. As with research on pregnant women, fetuses, and neonates (and unlike research on prisoners), the relevant clauses for children are both about research that is not otherwise approvable. In other words, research that satisfies other criteria, such as not involving greater than minimal risk (Code of Federal Regulations 2009, §404), can be approved without having to resort to either of the two clauses I wish to criticize in this article. The first objectionable clause gives guidelines for research that is “likely to yield generalizable knowledge about the subject’s disorder or condition” (Code of Federal Regulations 2009, §406), whereas the second gives guidelines for research that “presents an opportunity to understand, prevent, or alleviate a serious problem affecting the health or welfare of children” (Code of Federal Regulations 2009, §407).

Note that the second clause §407 is strictly analogous to clause §207 for pregnant women, fetuses, and neonates; indeed, it uses the exact same language except that the phrase “pregnant women, fetuses, or neonates” is replaced by “children.” Thus, the same point I made about generalizing §407’s language to benefits applies here too. Section 406 is more interesting, in that it gives various guidelines that are appropriate if other people with the same disorder or condition as the pediatric subject benefit from the research. The same point about generalizing applies here, too, though instead of generalizing opportunities to benefits, we generalize knowledge to benefits. As with opportunities to understand, prevent, or alleviate problems (the language of §207 and §407), the main point of gaining generalizable knowledge about the subject’s disorder or condition (the language of §406), when you cannot help the subject, is to benefit others with that disorder or condition.

THE ARGUMENT

To facilitate ease of expression, I introduce the expression *to-kind benefit*. A to-kind benefit of research is a benefit that goes to people of the same kind as the subjects of that

research. Vulnerabilities are kinds. For example, research on children with cystic fibrosis (CF) that cures cystic fibrosis offers a to-kind benefit, because its benefit (a cure of cystic fibrosis) helps others with the same disease as its subjects. In contrast, *not-to-kind* benefits of research do not necessarily go to people of the same kind as the subjects of that research. For example, research on prisoners that cures diabetes offers a not-to-kind benefit, because many beneficiaries of the research are not themselves prisoners.

Using this terminology, my thesis is that harmful research with to-kind benefits is no less wrong than harmful research with not-to-kind benefits is. This thesis follows straightforwardly from two assumptions. The first assumption is that an action that harms one person and benefits another is typically no less wrong if the harmed person and benefited person share traits. The second assumption is that the shared trait of a particular vulnerability is no exception to the first assumption. I defend each of these assumptions in turn.

The First Assumption

We do not normally think that my harming you is better justified if we both are left-handed than if only one of us is. The same thing is true for traits that are perhaps more intimately connected with a person's identity, for example, gender. If I may not harm a man in order to benefit a woman, then I also may not perpetrate that same harm in order to benefit another man instead. Now, nothing I have said so far contends that shared vulnerabilities are irrelevant—that is the point of my second assumption, which I examine shortly. The point here is simply that shared properties are typically insufficient to generate moral justificatory power.

Consider the trolley problem. If we must decide whether it is permissible or even obligatory to flip a switch that will send an out-of-control trolley down one track where it kills one, or let it continue on its way down another track where it kills five, it is irrelevant whether the six people share any traits in common, such as handedness or gender. Of course, the harmed subjects in research are used as means to generating the benefits to others, while the sacrificial victim in the trolley problem just happens to be at the end of the alternative track. However, things are no different if we move to the fat man variant of the trolley problem, wherein we must decide whether to push a fat man off a ledge and onto the track. If we do, he will fall and stop the trolley before it hurtles into the five, though he will die in the process. In this variant, if we push the fat man in order to save the five then we have used him as a means to saving their lives (Thomson 1985), but whether he shares any traits (handedness or gender, for example) with the five at the end of the track is again irrelevant to the moral question of whether we may push him, just as it is irrelevant to the moral question of whether we may flip the switch in the original trolley problem. We do not think, for example, that we may push the left-handed fat man if the five we are using him as a means to save are also left-

handed fat men, but not if they are right-handed thin women.

Note that my claim is not about what we may or may not do in any variant of the trolley problem; it is that whether victim and beneficiaries share properties typically is not relevant to those questions. Note also that my first assumption is compatible with various properties themselves having moral significance. For example, if species membership has moral significance, then maybe harmful research on humans is more difficult to justify than harmful research on pigs is. Even if so, however, being members of the same species does no further moral work. Harmful research on humans may be easier to justify if it benefits humans than if it benefits pigs, but if so this is because, according to our assumption, humans have higher moral status than pigs do, not because the beneficiaries are members of the same species as the harmed research subjects. After all, harmful research on pigs is easier to justify if it benefits people than if it merely benefits other pigs. The point of the first assumption is not that no trait has any moral significance; that would be absurd. Rather, it is that the mere sharing of traits typically does not making harming one in order to benefit the other any better.

The Second Assumption

My first assumption—that shared traits between the victims of a harm and its beneficiary typically do not make that harm any less wrong—might admit of exceptions. Consider, for example, nationality. Redistributive taxation from the rich to the poor within the same country seems better justified than redistributive taxation from the rich in one country to the poor in another country. Similarly, research that harms members of one country might seem better justified if the beneficiaries of that research are citizens of that same country. A plausible rationale for this exception to my first assumption appeals to the idea of a social contract: Citizens of the same country benefit from contracting with each other for mutual gain, but not so for citizens in different countries (though see Huemer 2013, chaps. 2 and 3).

This is not the place for a thorough examination of the moral or political significance of nationality; for our purposes it suffices to note that there is no similarly plausible rationale for the moral significance of shared vulnerabilities. Even if citizens of the same country benefit more from entering into a social contract with each other than with citizens of other countries, and even if this differential benefit can explain different political obligations, it is not similarly plausible that people with the same vulnerability benefit more from entering into a social contract with each other than with others, nor that this differential benefit can explain the moral significance of shared vulnerability in research.

I elaborate here with an example. If children with cystic fibrosis are vulnerable, then harms to one child with cystic fibrosis are not more justifiable just because the beneficiary of that harm also has cystic fibrosis. Again,

consider the trolley problem to make the issue concrete. Even if the trolley problem is difficult, it is no easier if all six people on the track are children with cystic fibrosis; that information is irrelevant. And, again, the fact that research subjects are used as means is also irrelevant, because the same points hold for the fat man variant of the trolley problem. In other words, we cannot justify a difference in our actions—pushing the fat man when they all have cystic fibrosis, not pushing when he does but they do not—by saying “at least the beneficiaries of our pushing him off the ledge have cystic fibrosis, just as he did.”

To be fair, the preceding discussion ignores a point I made earlier, that the vulnerability in question for to-kind benefits must be relevant to the benefit of the research. Having cystic fibrosis is clearly not relevant to death by being struck by a run-away trolley, but it might be relevant to respiratory failure. However, restricting the argument to relevant benefits does not affect its force. If harming you in order to benefit me is wrong, then such harm will not cease being wrong solely in virtue of the fact that the benefit to me is somehow related to our shared vulnerability. We can see that this is true by using the same technique I have already implicitly been using, the so-called bare difference strategy. As in James Rachels’s classic use of that strategy (Rachels 1975), we compare two cases that are as similar as possible in all respects except for the bare difference in question, here the presence or absence of shared vulnerability. Intuitively, that bare difference does not change our moral assessments of the two cases, and so is morally irrelevant in general. Thus, compare the following two cases:

Case 1: Killing to Cure CF

A 15-year-old has cystic fibrosis, and this disease will cut his life span to 30 years, but (for simplification) it will have no effect on the quality of his life. If you kill this teenager now you will of course deprive him of 15 years of life, namely, the 15 years he would otherwise get until he died of cystic fibrosis. Now, we could kill this teenager right now in order to cure five other teenagers of cystic fibrosis. They too are now 15 years old and will live to 30 if they have cystic fibrosis, but they will live to 80 if they are cured, with no change in quality of life.

Case 2: Killing to Cure Infection

Again, a 15-year-old has cystic fibrosis, and this disease will cut his life span to 30 years, but (for simplification) it will have no effect on the quality of his life. Again, if you kill this teenager now, you will of course deprive him of 15 years of life. In this case, however, we could kill the teenager now in order to cure *infections* that five other teenagers have. The original teenager does not share this infection, and it is not related to his cystic fibrosis, but he has some antibodies to it, and the only way to harvest these antibodies necessitates his death. These other five teens are also now 15 years old and will live to 30 if the infection persists, but they will live to 80 if they are cured of it, again with no change in quality of life.

In both cases we must decide whether to sacrifice one in order to save five. The only difference between the two cases is whether the benefits accrue to teens with the same disease as the harmed subject. In case 1, they all have cystic fibrosis, but in case 2 they have an unrelated infection that the harmed subject lacks. Intuitively, this difference is morally irrelevant; that fact has no bearing on what we may or may not do in the two cases.

Now, these examples differ from our original question in that they involve very severe wrongs, namely, killing. In contrast, our original question deals with a much more benign question, namely, whether we are allowed to recruit subjects to participate in a study to which they presumably consent. The disanalogy is useful, however, because it helps magnify any possible difference in intuition we might have about cases of wrongdoing done to benefit those who do, or do not, share a vulnerability with the harmed subjects. I have argued that there is no such difference, but the same conclusion applies if we reduce the severity of the harm in question. Consider, then, two further cases, which again are intended to be as similar as possible except for the one difference under examination, whether benefit is to-kind:

Case 3: Harmful Research to Alleviate CF

A 15-year-old has cystic fibrosis. If you conduct research on him you will cause him to experience considerable but temporary pain. The research conducted has some chance of extending the life span of other children with cystic fibrosis.

Case 4: Harmful Research to Alleviate Infection

Again, a 15-year-old has cystic fibrosis. If you conduct research on him you will cause him to experience considerable but temporary pain, to the same degree as in case 3. The research conducted has some chance of extending the life span of other children with an infection that is similar in its morbidity and mortality profile to cystic fibrosis, but distinct from it.

Again, the cases are similar in that we must decide whether to ask one person to sacrifice for the benefit of others. They differ in that the benefit in case 3 is to-kind, whereas the benefit in case 4 is not. Intuitively, this bare difference is morally irrelevant: We may ask the teenager in case 3 to participate if and only if we may ask the teenager in case 4 to participate; the fact that the subject in case 3 (but not in case 4) has the same illness as the targeted beneficiaries, and the fact that the benefit in case 3 (but not in case 4) is related to its subject’s vulnerability, are irrelevant.

Of course, other considerations are relevant; for example, the intensity and duration of the pain matters, as does the chance of extending life span. Similarly, whether we are exploiting the subject in order to reap the gains of research matters, as does whether the subject’s participation was essential in deriving the benefit of research. For

example, assume that we should not ask the subject to participate in case 4 if the benefit of research could have been derived without his participation. Then we should think the same in case 3: We should not ask its subject to participate if the benefit of that research could have been derived without his participation. (Note again that I am not making any claim as to what should actually be done in any of these cases; my claim is only that shared vulnerability is irrelevant to that question.) In contrast, the one difference between the two cases—the presence or absence of a shared vulnerability (which is related to the benefit received)—seems, on its own, irrelevant.

Scope

I have argued that, holding fixed all other relevant variables, whether the beneficiaries of research share some vulnerability with its subject is irrelevant to the moral question of whether we should conduct that research. At this point one might concede the soundness of this bare difference argument while suggesting that in the real world, whether the benefits of research are to-kind typically correlates with other morally relevant differences too. For example, maybe research with to-kind benefits is less likely to exploit its subjects or use them unnecessarily than research with not-to-kind benefits is. I address this objection shortly, but I want to conclude this section by saying a bit more about the scope of my argument.

First, CFR 45.46 does not sort various research proposals into the simple categories of “harmful” and “not harmful.” For research on prisoners, for example, the issue does not even arise, presumably because prisoners are autonomous adults, so that their informed decision about whether to participate trumps paternalistic weighing of harms versus benefits. (However, see Code of Federal Regulations 2009, §305a3, for a related clause about the commensurability of the research’s risks to risks that would be accepted by nonprisoner volunteers.) The closest CFR 45.46 comes to differentiating harmful from nonharmful research is in language in various places describing cases where the risks of research are greater than minimal, with no prospect of direct benefit (Code of Federal Regulations 2009, §204 and §404–406). These concepts are not equivalent, of course. Research with only minimal risks (e.g., blood draws) might still be harmful, as might be research with a small prospect of direct benefit yet certain harms (e.g., Phase 1 toxicity trials, interpreted as having a tiny chance of cure, but where the doses will be titrated up to toxic concentrations). Still, while there is a conceptual difference between my simpler category (harmful) and CFR 45.46’s more complex one (greater than minimal risk and no prospect of direct benefit), the simpler category of harmful research is broader and therefore of more interest than CFR 45.46’s category of greater than minimal risk with no prospect of direct benefit.

A related second point is that CFR 45.46 typically counts benefits toward research subjects as morally significant only if those benefits are direct (Code of Federal

Regulations 2009, §204d and §405). As an example, a direct benefit of cystic fibrosis research is a cure for cystic fibrosis, whereas an indirect benefit of the same research might be better ancillary medical care that the subjects receive while participating in the study. Whether the Code of Federal Regulations is right to distinguish between benefits in this way is controversial (Wertheimer 2011; Friedman, Robbins, and Wendler 2012), but because the benefits at stake in this article can easily be limited to direct benefits (e.g., curing cystic fibrosis rather than getting better ancillary care), we may safely ignore the controversy.

Third, my argument is compatible with several other plausible if controversial theses about the moral justifiability of harming one in order to benefit another. For example, it is compatible with the thesis that the magnitude of benefits to others, relative to the harm to the sacrificing subject, is morally important. If we could cure cystic fibrosis by sacrificing one unlucky cystic fibrosis sufferer, probably we should do so. Of course, it is equally true that if we could cure cystic fibrosis by sacrificing one unlucky diabetic, we should probably do that. Thus, magnitude of benefit relative to harm matters, while shared vulnerability does not.

My argument is also compatible with the thesis that if the harmed subject is willing to sacrifice in order to save others with the same condition that he has, but unwilling to sacrifice to save others a similar magnitude of harm from some other condition that he lacks, then his wishes should be respected. A person with cystic fibrosis might identify more closely with people who share his disease, for example, than with people who have some other disease. That is fine, but the issue at stake is whether this possibility can justify precluding researchers from even asking for consent to participate in research that benefits people with some other disease. It might be imprudent, in the sense of being a waste of time, for researchers to ask for a candidate subject’s consent if they suspect that he will not identify closely with the intended beneficiaries of research, but that is not the same as suggesting that the researchers should be forbidden as a matter of federal policy from asking for consent in such cases.

My argument is also compatible with the thesis that the welfare of the target beneficiaries, relative to other members of society, is morally relevant. For example, if the target beneficiaries in the two “Harmful Research” cases are among the worst off people in society, maybe we should prioritize their welfare over that of the subject and therefore ask for his participation. But this is true regardless of whether the benefit to the five is reducing the effect of cystic fibrosis, a condition the subject has as well, or reducing the effect of an infection, a condition the subject lacks.

Another way to phrase these points about compatibility is to shift the burden of proof: We have no good reason to think that sharing a vulnerability should be relevant to the morality of harming one person to benefit someone else; that claim is counterintuitive and does not follow from any more general ethical principle or theory. In contrast, we have good reason to think that the relative

magnitude of benefit to harm can be morally relevant. We likewise have good reason to think that the target's differential consent can be morally relevant. And, finally, we have good reason to think that the welfare of the beneficiary, relative to that of others in society, can be morally relevant. In contrast, whether the target beneficiary stands to gain from something connected to some vulnerability that might be shared by the harmed subject seems, absent further argument, irrelevant.

Finally, my conclusion should not be interpreted as suggesting that we should never conduct research that harms its subjects. I acknowledge of course that such harmful research may still be morally permissible, even praiseworthy or obligatory. As an analogy, abortion clearly harms a fetus, and yet it might be morally permissible, perhaps because women should have rights over their bodies. Similarly, even if research harms fetuses, such harm may be trumped by other concerns. My conclusion is in this sense quite cautious. I am not claiming that various kinds of harmful research should not or may not be conducted; I am saying only that vulnerability-sharing between subject and beneficiary is irrelevant to the moral question of whether harmful research should or may be conducted.

VULNERABILITY AS PROXY

An obvious response to my argument is that, while sharing a relevant vulnerability is, strictly speaking, morally irrelevant, it is a useful proxy for something that is morally relevant. There are three ways that this response might continue, but all of them fail.

Prioritarianism

Prioritarianism is the view that we should give moral priority to the worst-off members of society. Might vulnerability be a proxy for being badly off? I concede that vulnerable people may be worse off in society than the nonvulnerable are. However, this point cannot help support CFR 45.46, for several reasons. First, even if the goal of CFR 45.46's restrictions is to prioritize the vulnerable, permitting harmful research only if its benefits are to-kind is too restrictive. After all, if we want to prioritize benefits to children, we should be willing to conduct research on prisoners in order to benefit those children. This is not allowed under CFR 45.46, however; harmful research on prisoners must benefit prisoners, even if we would like to prioritize some other vulnerable group, such as children, instead. In other words, prioritarianism does not distinguish between different vulnerabilities (e.g., being a prisoner or being a sick child) that make its targets worthy of priority, whereas CFR 45.46 does. Therefore, prioritarianism cannot justify CFR 45.46's restriction that harmful research must have to-kind benefits, in which case CFR 45.46's restrictions are not an adequate proxy for prioritarianism.

Also, it is hard to see how being pregnant or being a child thereby makes one worse off than others in society.

A child, for example, may be especially vulnerable to medicines that have unknown effects on children, but that is not the same as saying that the child is, all things considered, worse off than adults. Likewise, a pregnant woman may be especially vulnerable to unknown interactions between medicines and her fetus, but that is not the same as to say that she is, all things considered, worse off than people who are not pregnant.

Still, there may be a correlation between being vulnerable and being badly off. However, even if vulnerable people tend to be worse off than the nonvulnerable, vulnerable people are not necessarily the worst off members of society. There are rich children, for example. Indeed, if CFR 45.46's concern were really about the worst off, then it has neglected the most obvious trait that diminishes welfare: poverty. In other words, CFR 45.46's restriction to research with to-kind benefits is a poor fit with prioritarianism, first because the vulnerable are not necessarily badly off, let alone among the worst off members of society, and second because there is another group that is clearly badly off—the poor—that the Code of Federal Regulations ignores entirely.

Exploitation

According to this defense of CFR 45.46, restricting research more heavily when its benefits are not-to-kind is a proxy for reducing the chance of exploitation. For example, if harmful research on a subject with cystic fibrosis must face stricter ethical scrutiny when such research benefits people who lack cystic fibrosis, then we will be less likely to exploit subjects with cystic fibrosis by conducting research on them.

Now, there is some controversy about what exploitation is and when, if ever, it is wrong (Wertheimer and Zwolinski 2013), but regardless of the right answer in those controversies the appeal to exploitation is mistaken. If we really care about protecting vulnerable subjects from possible exploitation, we should not let our guard down just because the research in question benefits others with the same vulnerability as the subject. That is true because to-kind benefits neither eliminate nor even reduce the possibility of exploitation in research. For example, whether we judge a case of research on subjects with cystic fibrosis as exploitative depends on how those subjects are treated, not on who the beneficiaries of that research are. Analogously, whether sweatshop labor exploits Southeast Asian workers depends on how those workers are treated, not on whether the beneficiaries of that exploitation—the factory owners and the consumers who purchase sweatshop items at low prices—are themselves Southeast Asians. We can concede that a factory owner who sells his products to Southeast Asians at low prices benefits the general community of Southeast Asians, even while exploiting his workers in particular to do so. Similarly, a researcher might benefit the general community of people with cystic fibrosis even while exploiting her research subjects to do so.

Indeed, if anything, CFR 45.46 is susceptible to a further accusation here, that it is more liable to license exploitative research just because the beneficiaries of said research share some vulnerability with its harmed subjects (see Iltis 2011). That is because, as I have just said, whether research is exploitative does not depend on whether its beneficiaries and subjects share vulnerabilities. Therefore, if we are more lenient in allowing research when subjects and beneficiaries share some vulnerability, we are also more likely to allow exploitative research in those cases.

There is a possible response to what I have just argued: Harming a vulnerable subject in order to benefit someone else is exploitative if the beneficiary is better off than the subject is, but not otherwise. Research beneficiaries who share a relevant vulnerability with its harmed subjects are, at least by that standard, no better off than those subjects, and such research is therefore not exploitative.

This response will not do, however. First, sharing the same vulnerability is at best very loosely correlated with relative welfare, all things considered. For example, one person with cystic fibrosis may be much worse off than another, in multiple ways, such as income, education and employment prospects, and so on, that are independent of cystic fibrosis. The same applies for two prisoners, and for two pregnant women, and so on. Even if we want to restrict our welfarist concerns to medical welfare, so that sharing the same disorder is more significant (because nonmedical aspects of welfare are ignored), two children can have very different levels of health, as can two people with cystic fibrosis. Thus, if we really care about the subjects of harmful research being no worse off than its intended beneficiaries, we should be explicit about this and write it into the Code of Federal Regulations.

A second, related point is that restricting welfarist concerns to the presence or absence of vulnerabilities would be as morally arbitrary as the distinction between research with to-kind benefits versus not-to-kind benefits, and the whole point of this response was to argue that what seemed like an arbitrary moral distinction (whether benefits are to-kind) was in fact justified because it turns out to be a useful proxy for a moral distinction that is not arbitrary (relative welfare levels). If the only way to defend the latter moral distinction is by positing further moral arbitrariness (that the only kind of relative welfare that matters is the vulnerability itself) then this rejoinder has not gained any ground.

Third, and perhaps most importantly, no plausible account of exploitation entails that what would otherwise be harmful exploitation is not exploitative if the beneficiary of the alleged exploitation is as badly off as the harmed target of the alleged exploitation. Analyses of exploitation can be divided into roughly two types. One type argues, in a roughly Kantian vein, that exploitation occurs when you fail to treat people with proper respect (Sample 2003). On this sort of view, whether the beneficiary of an exploitative act is as badly off as the exploited person is irrelevant to the charge of exploitation, because relative welfare levels are not relevant to the charge of failing to treat with proper

respect. For example, enslaving someone fails to treat that person with proper respect even if the slave owner is as badly off as the slave is.

The other sort of account of exploitation argues, in a fairness-based vein, that exploitation is unfairness in the sense of insufficient benefit to the exploited party (Wertheimer 1996). For example, if the sweatshop laborer's salary is too low then it is possible that he is being exploited. This sort of analysis is also unhelpful in the present discussion, however, because it focuses on a different comparison than the one being proposed. The analysis of exploitation as insufficient benefit compares what the subject's benefit is to what it should be, but the proposal we are currently considering compares the starting welfare levels of both parties. Again, a poor factory owner can exploit his laborers via low wages just as easily as a rich factory owner can.

Consider a final variant of the "vulnerability as proxy for exploitation" defense of CFR 45.46, according to which researchers whose studies benefit a vulnerable group are likely to care about that group, which in turn reduces the chance that those researchers will exploit others in that group. The problem with this variant defense is that it is vulnerable to a dilemma: Its empirical claims are either false or insufficient to prove the desired conclusion, depending on how they are interpreted. This variant defense asserts two correlations, first between benefiting members of a vulnerable population and caring about that population and second between caring about a vulnerable population and not exploiting members of that population. The problem is that tight correlations are needed for the defense to succeed, but only loose correlations are plausible.

Tight correlations are implausible: There is no guarantee that a researcher conducting research that benefits a vulnerable population thereby cares about said population. For example, a particular researcher might conduct research intended to cure cystic fibrosis without caring at all about people with that disease; she might care only about her own career. Along the same lines, a researcher who cares about a vulnerable population might still exploit some members of that population in order to benefit others in the population. Indeed, a utilitarian might say that we are obligated to exploit the vulnerable, when doing so maximizes aggregate utility. Thus, for example, a researcher who cares about children with cystic fibrosis might still exploit some such children in order to benefit others with the disease.

Loose correlations are insufficient: They do not protect adequately against exploitation. As an analogy, we do not normally think that citing loose correlations about caring is sufficient to preclude exploitation of factory workers: A factory owner who intends to sell low-cost sport shoes to Southeast Asians (thereby benefiting them) might still exploit the Southeast Asians under his employ, and suggesting that there is a loose correlation between benefiting Southeast Asians and caring about them, and between caring about Southeast Asians and not exploiting them, is not

a sufficient protection against that exploitation. A better defense against exploitation, both for vulnerable research subjects and for factory workers, would be to operationalize just what exploitation means (either as the Kantian injunction against treating as mere means or as the injunction to ensure that benefits are fair) and then have explicit regulations guarding against it.

Necessary Sacrifice

According to this last reply, the point of keeping track of whether research's beneficiaries share a vulnerability with its harmed subjects is to ensure that we conduct harmful research on vulnerable subjects only when absolutely necessary (see Jonas 1969, 241). The guiding idea is that if we could get the same benefits of the research by conducting it on nonvulnerable subjects then we should do that instead of recruiting subjects who are already vulnerable.

The guiding idea—that we should conduct research on the vulnerable only when necessary—is laudable, but it does not support restricting research on the vulnerable to cases with to-kind benefits. First, and perhaps most obviously, ensuring that the beneficiaries of research share the same vulnerability as its subjects does not guarantee the participation of vulnerable subjects was essential. For example, even if harmful research conducted on children with cystic fibrosis is intended to benefit others with cystic fibrosis, it is possible that the knowledge gained from such research could have been gained just as easily in other ways, for example, by switching to a study whose subjects are healthy adults.

Now, one might respond to what I just said by weakening the original idea. In particular, one might say that restricting to research with to-kind benefits makes more likely, though does not guarantee, that harming vulnerable subjects was the only way to accrue that research's benefits. However, while this weaker claim about increasing likelihood is certainly true in one sort of scenario, it is not true in another, more commonly encountered scenario. Suppose first that we fix the assumption that we are conducting research on children with cystic fibrosis. Under this fixed assumption, we reduce the chance that we are calling upon that population frivolously if the intended benefit of the research is related to cystic fibrosis rather than to diabetes. If the intended goal of research on children with cystic fibrosis is to alleviate diabetes, then one can rightly question whether the researchers really needed children with cystic fibrosis in the research. But the contributions of subjects with cystic fibrosis are often essential to the outcome of research if the goal of that research is to alleviate cystic fibrosis.

On the other hand, suppose we fix that the intended goal of the research is to benefit children with cystic fibrosis. Then we may ask whether, given this fixed goal, one must recruit children with cystic fibrosis to conduct the research. Here the answer is often “no,” for there will be many scenarios where we can conduct research on nonvulnerable subjects in order to benefit vulnerable populations.

Phase 1 toxicity trials are examples; in such cases we could recruit children with cystic fibrosis, but we can often get sufficient (and, indeed, potentially better) results by recruiting healthy people instead. Thus, if the goal of restricting research on the vulnerable to that with to-kind benefits were to make more likely (or guarantee) that we conduct research on the vulnerable only when necessary, that restriction would frequently fail to achieve its goal, because it would not preclude the use of vulnerable populations in situations where we could instead recruit nonvulnerable subjects to benefit the same vulnerable group.

In the first scenario I discussed, I considered cases where we fix our research subjects and their particular vulnerability, and we then ask whom the research might benefit. In the second scenario I discussed, I considered cases where we instead fix the goal of research and the particular vulnerability it intends to alleviate, and we then ask whom we might enroll as participants in that research. The second scenario more closely mirrors the actual order of operations for most research. Rarely will a researcher first settle on the fact that she wants to conduct research on children with cystic fibrosis and then ask who might benefit from said research. More common by far, I suspect, are scenarios where the researcher knows what the intended finding or goal of the research is, and then asks who might be appropriate subjects for the study. In these more common cases, restricting research to that with to-kind benefits does not sufficiently reduce the chance that vulnerable subjects will be asked to participate needlessly.

There is a second reason that to-kind benefits are a bad proxy for the necessity of sacrifice: The proxy is no clearer or easier to follow than the thing for which it is the proxy. For example, if we really care about ensuring that prisoners in some sense had to participate, in order to reap the research's benefits, we should just write that into CFR 45.46 explicitly, instead of resorting to the proxy of distinguishing between research whose benefits are to-kind versus not-to-kind. After all, CFR 45.46 already uses exactly this sort of clause—requiring that the use of certain subjects is necessary in order to achieve the gains of the research—in other areas. For example, in §204b, CFR 45.46 states that one necessary condition for permissible research on fetuses is that “the purpose of the research is the development of important biomedical knowledge *which cannot be obtained by any other means*” (Code of Federal Regulations 2009, §204b, my emphasis). Thus, the presence of to-kind benefits is not a useful proxy for the necessity of research, because it would be easier and clearer just to be explicit about the moral importance of the necessity of using vulnerable subjects, rather than to try to get at that idea through the proxy.

CONCLUSION

In this article, I have argued, against CFR 45.46, that we should not be more permissive of harmful human subjects research when the beneficiaries of said research share important vulnerabilities with its harmed subjects. This

conclusion applies for every vulnerable population explicitly discussed by CFR 45.46: pregnant women, fetuses, neonates, prisoners, and children, including children with specific conditions or disorders. What sort of policies should we implement instead, then? After all, surely vulnerable populations demand special protections. The answer to this question will depend on the particular vulnerabilities in question; no blanket statements will be possible (see Levine, Faden, et al. 2004).

Of course, my thesis is compatible with the laudable goal of protecting the vulnerable from harmful research. The point I have been defending in this article is that, when trying to articulate just when we can make exceptions to the general rule of protecting the vulnerable from research, we should not rely on the distinction between to-kind and not-to-kind benefits. Maybe instead we should be content with the conclusion that the morality of research should be judged on many merits, such as five I discussed in this article: (1) relative magnitudes of benefits and harms, (2) whether the subjects consent, (3) whether the beneficiaries are among the worst off in society, (4) whether the subjects are being exploited, and (5) whether the participation of vulnerable populations was necessary for the research goals. (see also Emanuel, Wendler, and Grady 2000). The thesis of this article is that whether the beneficiaries of research share particular vulnerabilities with its harmed subjects is not one such merit to be listed along with these others, and it is not even a good proxy for any of them either. ■

ACKNOWLEDGMENTS

The author would like to thank participants at a University of Colorado at Boulder Center for Values and Social Policy talk for their helpful discussion. Special thanks are owed to David Boonin and Adam Hosein for their helpful comments on early drafts of this article.

REFERENCES

- Code of Federal Regulations. 2009. Title 45, Public Welfare, Department of Health and Human Services, National Institutes of Health, Office for Human Research Protections, Part 46, Protection of human subjects. Bethesda, MD. Available at: <http://www.hhs.gov/ohrp/humansubjects/guidance/45cfr46.html> (accessed May 7, 2014).
- Emanuel, E. J., D. Wendler, and C. Grady. 2000. What makes clinical research ethical? *Journal of the American Medical Association* 283 (20): 2701–2711.
- Friedman, A., E. Robbins, and D. Wendler. 2012. Which benefits of research participation count as ‘direct’? *Bioethics* 26(2): 60–67.
- Huemer, M. 2013. *The problem of political authority: An examination of the right to coerce and the duty to obey*. New York, NY: Palgrave Macmillan.
- Iltis, A. 2011. Justice, fairness, and membership in a class: Conceptual confusions and moral puzzles in the regulation of human subjects research. *Journal of Law, Medicine & Ethics* 39(3): 488–501.
- Jonas, H. 1969. Philosophical reflections on experimenting with human subjects. *Daedalus* 98(2): 219–247.
- Levine, C., R. Faden, C. Grady, et al. 2004. The limitations of “vulnerability” as a protection for human research participants. *American Journal of Bioethics* 4(3): 44–49.
- National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1979. *Belmont report: Ethical principles and guidelines for the protection of human subjects research*. Bethesda, MD. Available at: <http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.html> (accessed May 7, 2014).
- Rachels, J. 1975. Active and passive euthanasia. *New England Journal of Medicine* 292(2): 78–80.
- Sample, R. 2003. *Exploitation, What it is and why it is wrong*. Lanham, MD: Rowman and Littlefield.
- Thomson, J. J. 1985. The trolley problem. *Yale Law Journal* 94(6): 1395–1414.
- Wertheimer, A. 1996. *Exploitation*. Princeton, NJ: Princeton University Press.
- Wertheimer, A. 2011. Is payment a benefit? *Bioethics* 27(2): 105–116.
- Wertheimer, A., M. Zwolinski. 2013. Exploitation. *The Stanford Encyclopedia of Philosophy* (Spring 2013 edition). Available at: <http://plato.stanford.edu/entries/exploitation/> (accessed May 7, 2014).