

ELIZABETH HARMAN

HARMING AS CAUSING HARM

1. INTRODUCTION

In this paper, I will offer a solution to the non-identity problem and defend that solution. The non-identity problem arises because some actions appear to be wrong, and they appear to be wrong in virtue of harming certain people, but those people would not have existed if the actions had not been performed, and those people have lives that are worth living. Such actions are puzzling because they do not make these people worse off than they otherwise would have been; but plausibly, one harms someone only if one makes her worse off. A solution to the non-identity problem would both explain why the actions are wrong and vindicate the appearance that the actions are wrong in virtue of harming the relevant people.¹

Consider the following case. A woman has a temporary condition such that if she conceives now, her child will be blind. However, if she waits for three months, she can conceive a child who will not be blind. Clearly, this woman should not conceive now. It would be wrong to conceive now. Furthermore, it seems that if she does conceive now she is *harming* her future child, and *that* is why her action is wrong. However, it can be argued that she does not harm her child and thus does not act wrongly. Plausibly, one harms someone only if one makes him worse off than he would otherwise have been. But this woman's act of conceiving now does not make her child worse off than he would otherwise have been. Had she not performed this action, he would not exist. Furthermore, he has a life worth living. Either he is better off than he

would otherwise have been (because existence can be compared to non-existence, and it is better to have a life worth living than not to exist) or he is not worse off because he is neither worse off nor better off (because existence cannot be compared to non-existence). Either way, he is not worse off than he would have been had she not performed this action. Thus, it seems that she does not harm him. Furthermore, it seems that the only candidate explanation of her action's wrongness is that it harms; so it seems her action is not wrong.

Another case involves the disposal of nuclear waste. We are trying to decide whether to enact a lax nuclear waste disposal policy. The differences between our current policy and the lax policy are considerable: they involve which disposal plants are built and which jobs are available, and they require different things of ordinary people in their daily lives (such as recycling policies). If we do enact the lax policy, this will be more convenient for us; but we know that several hundred years from now, many people will become sick from the poorly contained waste: many people will suffer from cancer and die young. However, we are also able to predict that *who* exists in the future will differ depending on whether we enact the policy: who has which jobs, who lives where, who meets whom—all of these things will differ in large and small ways and ultimately affect who has children together, or at least when people conceive. Enacting the policy seems to be wrong, and it seems to be wrong because it harms future people, but it does not make them worse off than they otherwise would have been.

The above two cases are cases about which there is widespread intuitive agreement. Pre-theoretically, the actions seem clearly to be wrong. Solving the puzzle

would involve vindicating our intuitive beliefs that the actions are wrong, as well as our intuitive beliefs about *why* they are wrong.²

There is another kind of case that raises the non-identity problem, but about which there is not widespread agreement. Indeed, there is real intuitive disagreement about these cases. One example is the following case. Suppose that a woman has a permanent condition such that if she ever conceives a child, that child will be deaf. Is it morally permissible for her to conceive? Pre-theoretically, people disagree. Another example is the following. A man and a woman are married, in love, and want to have children together; but they are slaves in the American South, in 1800. If they conceive, their child will be a slave too. Is it morally permissible for them to conceive? Again, pre-theoretically, people disagree.

A solution to the non-identity problem would be better if it could explain why for some cases, intuitions are uniform, while for other cases—those I will call “hard cases”—there is disagreement. I will develop a solution that can explain this.

2. HARMING AS CAUSING HARM

The view I will offer, which provides a solution to the non-identity problem, begins with three claims about harm.³ The first is a sufficient condition on harming:

- (1) One harms someone if one causes him pain, mental or physical discomfort, disease, deformity, disability, or death.

(The slogan version of claim (1) is: harming is causing harm. That is, an action is a harming action just in case it causes an effect of harm.) More generally, the view is that an action harms someone if it causes the person to be in a bad state. Bad states are understood as states that are in themselves bad, not bad because they are worse than the

state the person would otherwise have been in. (If one wants a further account of a bad state, I am willing to offer one: bad states are those states that are worse in some way than the normal healthy state for a member of one's species. See my 2004.)

The most powerful intuitive claim in favor of the idea that one cannot harm someone unless one makes him worse off is the following: surely a doctor does not harm me when he gives me a life-saving operation. Nevertheless he causes me pain and he physically injures me, as well as doing whatever saves my life. On my view, the doctor does harm me. However, the doctor *permissibly* harms me, because the doctor prevents greater harm to me than he causes. Furthermore, seeing that the doctor harms me enables us to see why surgery is a last resort and is not to be undertaken lightly—we must be quite sure the harm is justified before proceeding.

The import of the claim that someone harms someone else is given by this claim:

(2) An action that harms someone thereby has a strong moral reason against it.

Furthermore, while the prevention of worse harm to someone can justify harming her, the mere fact that a harming action also *benefits*—that is, provides positive good things to the harmed person—does not justify the harm. On my view:

(3) The mere fact that a harming action also benefits the person harmed, and benefits her more than it harms her, is insufficient to justify the harm.⁴

Besides these three claims about harm, my view also makes a fourth claim about one kind of harm: harm to people who do not exist independently of the harming action. Some actions harm people who exist regardless of whether the action is performed. Other actions (including the actions in non-identity cases) harm people who would not

have existed if those actions had not been performed.⁵ The fourth claim elaborates a particular way that reasons against harm are *strong* and so harming is not easily justified.

- (4) If an action harms someone who does not independently exist, then the fact that the action also benefits the person harmed, and benefits him more than it harms him, is *ineligible* to justify the harm if *failing to perform the action* would similarly benefit someone.

The intuitive thought behind claim (4) is the following. The fact that an action harms someone provides a strong moral reason against acting; it is a reason that tells in favor of *refraining* as opposed to *performing the action*. If the action also benefits the harmed person, but performing the action is not the only way to provide such benefits—indeed, refraining from performing the action would provide similar benefits to someone—then considerations of benefit simply do not tell *in favor of acting* as opposed to *refraining from acting*. On my view, considerations of benefit are therefore “ineligible” to justify the harm in that they never justify the harm in such cases. (In Section 4 I discuss why claim (4) is restricted to harms to people who do not independently exist.)

I assume that counterfactual dependence is sufficient for causation: I assume that if an action is such that if it had not been performed, a particular event would not have occurred, then the action causes the event. While counterfactual dependence is clearly not necessary for causation, as cases of preemption illustrate, it is commonly taken to be sufficient for causation.

3. SOLVING THE NON-IDENTITY PROBLEM

A “non-identity case” is a case in which an action affects whether some people exist, and it appears (either to everyone, or to some people) that the action would be

wrong in virtue of harming these people, although they have lives that are worth living. I will explain how my view solves the non-identity problem by discussing four different types of non-identity cases:

- (i) same-number cases: cases in which the same number of people will exist in the future regardless of which way the agent acts
- (ii) different-number cases, more people if the action is not performed: cases in which different numbers of people will exist depending on whether the agent performs the action, and the agent knows that more people will exist if she does not perform the action
- (iii) different-number cases, unknown which scenario has more people: cases in which different numbers of people will exist depending on whether the agent performs the action, and the agent does not know whether more people will exist if she performs the action (and does not know whether more people will exist if she does not)
- (iv) different-number cases, more people if the action is performed: cases in which different numbers of people will exist depending on whether the agent performs the action, and the agent knows that more people will exist if she performs the action

I will begin with same-number cases. One example of a same-number case is the temporary condition case I described at the beginning of this paper. A woman is deciding whether to conceive now, while she has a temporary condition that would cause her child to be blind. She knows that if she does not perform this action, she will conceive in a few months, having a child who will not be blind. Another example is a case in which a

fourteen-year-old girl is deciding whether to conceive now. If she conceives now, her child will face many difficulties because she is not mature enough to be a responsible mother. If she does not conceive now, she knows that she will have a child later, when she is more mature. (I stipulate that the women in these two cases know they will conceive later if they do not conceive now.)

Each of these procreative actions causes the created child to be in a bad state—the woman in the temporary condition causes her child to be blind; the teenager causes her child to suffer the various hardships involved in having an immature mother. That each action *causes* these effects follows simply from the fact that had the action not been performed, the effects would not have occurred. Because each action harms the created child, there is a strong reason against each action. Now we must ask whether the harm is justified.

While some people deny that there is any harm in non-identity cases, claiming that an action cannot harm someone if it does not make him worse off, others grant that non-identity actions harm but claim that non-identity actions do not “all things considered” harm the created persons, because they also more greatly benefit the persons than they harm them.

On my view, non-identity actions both harm and benefit the created persons. They do benefit the created persons *more* than they harm them, in that the created persons’ lives are overall worth living: the benefits compensate for the harms in this sense. However, it is a further question whether the fact that the action benefits can *justify* the action’s harming. In same-number cases, there are strong reasons against the actions in virtue of the fact that the actions harm the created persons. Can the fact that

the actions also *benefit* these people justify the harming? On my view, this fact does not function simply to *cancel* the reasons against harm, because on my view the *mere fact* that a harming action also benefits the person harmed, and benefits him more than it harms him, is insufficient to justify the harm. Can this fact nevertheless justify the harming? It cannot, because it does not function as a reason to perform the action *rather than* to refrain from performing it. The reason against harming is a reason to refrain from acting: if she refrains, the agent won't harm anyone in the way she will if she acts. However, the action's benefiting is not like this. If the agent waits, she will have a different child later. Thus, whether she performs the action or refrains from performing it, she will benefit *someone* in the same way. So the fact that the action benefits provides no reason to act rather than refrain from acting, and on my view it does *no work at all* to justify the action's harming.

I turn now to different-number cases in which more people will exist if the action is not performed than if it is performed. I do not think cases like this are normally discussed in connection with the non-identity problem—indeed, they are somewhat artificial—but I will discuss them because they are relevant to the next category of case. Here is an example of such a case. A woman has a temporary condition such that if she conceives now, she will have a child who will be deaf and she will also become infertile after this pregnancy; that is, the temporary condition will result in both deafness of the child she conceives now and future infertility if she conceives now. However, if she waits, she will go on to have two children. (She and her partner have a slight but considered preference for two children rather than one, and they have every reason to

believe that they will be able to have two if they wait for the temporary condition to clear up.)

In this case, as for the prior type of case, there is a strong reason against the action because the action *harms* the child by causing his deafness. The action also *benefits* the child by causing all the good aspects of the child's life. Can the fact that the action benefits justify the action's harming? It cannot, because, as before, the benefits do not provide a reason to conceive now *rather than not*; indeed, in this case, the agent will benefit *more* people if she does not perform the action.

Thus, as for same-number cases, in this type of different number case, there is a strong reason against performing the action *rather than not performing it*, which is that it harms in a particular way, and there is no reason to perform it *rather than not* available to justify the harming. Thus, both types of actions are wrong in virtue of harming.

I turn now to different-number cases in which the agents do not know whether more people will exist if the action is not performed than if it is performed (and do not know whether fewer people will exist if the action is not performed than if it is performed). Cases of this type are commonly discussed in the literature on the non-identity problem, and while it is often noted that these are different-number cases, it is not generally noted that they are cases in which the agents *do not know* which outcome would have more people. I think this fact is very important, as will be clear in my discussion of these cases. An example of a case of this type is the nuclear waste disposal case. In this case, the agent is able to see that it's very likely that *who* exists in the future will be affected by whether the lax nuclear waste policy is put in place. By the same mechanism, it's very likely that *how many people* exist in the future will be affected, but

it is not at all clear which possibility would end up having more people in it. In this case, the agent should think it's overwhelmingly likely that how many people will exist in the future depends on whether the action is performed: she should have a very low degree of belief that this is a same number case. (While it's possible that it's a same number case, it would be a surprising coincidence if it were.) But she has no particular reason to favor the belief that *more* people will exist if she performs the action over the belief that *fewer* people will exist if she performs the action, and vice versa. So she should have a degree of belief of roughly 1/2 that this is a different-number case in which *fewer* people will exist if the action is performed than would exist if it is not performed. And she should have a degree of belief of roughly 1/2 that this is a different-number case in which *more* people will exist if the action is performed than if it is not performed.

In general, the following claim is true: if (a) an agent should have a degree of belief of roughly 1/2 that her situation is a particular way, and (b) if her situation were that way, then her action would be wrong, then (c) she should not act: she should do the morally cautious thing and refrain from acting.⁶ This claim applies to the case we have been discussing. In this case, the agent should have a degree of belief of roughly 1/2 that her action would be a non-identity action such that *fewer people* would exist if it is performed than if it is not performed. Such actions are wrong (for the reasons I gave above), so she should not act. To summarize: in different-number cases in which the agent does not know whether more people would exist if she performed her action (and does not know whether fewer people would exist if she performed her action), the actions are wrong *because* given what the agent knows about her situation, it is all too likely that her action would be wrong, and thus she should refrain from acting.⁷

Finally, I turn to discussing cases in which the agent knows that different numbers of people will exist depending on whether she performs the action, and she knows that more people will exist if she performs the action. Cases of this type are typically discussed in the literature on the non-identity problem; however, these cases are the *hard cases* I mentioned in the introduction. These are cases about which there is widespread pre-theoretical disagreement. For example, some people think it is permissible to procreate if one has a permanent condition such that *whenever* one procreates, one will have a deaf child; some people think this is impermissible. Some people think it is permissible to procreate if one is a slave whose child will be a slave; others think it is not.

In each of these cases, there is a strong reason against the action because it harms the created child by causing all the bad effects in the child's life—and in particular by causing the unusual bad effects (deafness in one case, the hardships of slavery in the other). Each action also substantially benefits the created child, by causing all the good things in his/her life. In these cases, it is not true that the agent will similarly benefit someone else if she refrains from acting, so the benefits do tell in favor of acting *rather than refraining from acting*. Thus, on the view I have developed so far, it is open whether the benefits justify the harming.

Note that cases of this type are very similar to cases of ordinary procreation. When deciding whether to procreate at all, one is deciding whether to perform an action which will affect *how many* people exist in the future and which will *harm* the created person by causing the bad aspects of his life. But cases of ordinary procreation are not “non-identity cases” as I have defined them, because they are not cases of actions that

seem to many people to be intuitively wrong. Rather, ordinary procreation is intuitively permissible. Cases of ordinary procreation are not *hard cases*.

I will now flesh out my view and show that it explains why ordinary procreation is permissible, but why procreating in the “hard cases” *may not be* permissible. On my view, an action that harms someone thereby has a strong reason against it. The mere fact that the action also benefits the people it harms—and benefits them more than it harms them—is insufficient to justify the action. Indeed, the fact that the action benefits is only eligible as a *possible justification* if the agent would not provide similar benefits to someone if he failed to perform the action—that is, if consideration of benefits tells in favor of acting *as opposed to refraining*. Furthermore, it is crucial that the benefits be benefits *to* the persons who are harmed by the action. One cannot in general justify severely harming one person by the fact that one will also provide positive benefits to other people.⁸ That the action benefits the very people whom it harms (and benefits them more than it harms them) is a *necessary* condition of the action’s being justified.

If this eligibility condition is met (refraining does not benefit in similar ways) and the necessary condition is met (the benefits are to the person who is harmed), then an action that harms someone *may* be justified. Whether it is justified is a matter of how serious the harm is, how substantial the benefits are, and perhaps how strong the other reasons for the action are (if there are any). My view is thus that whether harming is justified is a matter of threshold crossing. There is more than one threshold. Some harms are so bad that they are impossible to justify; so the harm must be below a certain threshold. And depending on how bad the harm is, there will be a corresponding threshold that the benefits must cross to justify the harming. (On a variant of my view,

other reasons for action may be relevant as well. The fact that a woman has a permanent condition such that whenever she conceives, she will have a deaf child, means that her only chance to have a child involves conceiving in these circumstances; there is a considerable sacrifice for her if she refrains. I do not believe that this consideration can play a role in justifying the harm of procreation, though there is something natural to thinking it can; I discuss this more below.)

On my view, while ordinary procreation involves harming the created person, the kind of harming involved is not so substantial that it cannot be justified. The benefits involved in ordinary procreation are so substantial that they do justify harming the person in this way. (It doesn't seem that *other reasons* to procreate play a role in making ordinary procreation permissible; it would be odd, for example, to say that having one's first child is more justified than later children because whether one has any children greatly affects one's life but whether one has more children doesn't greatly affect whether one has a good life—though of course it affects what one's life is like and the way it is good.)

When we think about ordinary procreation, it is natural to have thoughts along the lines I am expressing. It is natural, in imagining conceiving, to worry about all the hardships and bad things one's child's life will involve. It is natural—and on my view, correct—to think that these things will be compensated for by all the good things in one's child's life.⁹

Furthermore, on my view, there are some harms such that *nothing* could justify causing them. Suppose a woman has a permanent condition such that if she ever conceived a child, her child would have a disease that would at some point cause a year

of total agony, akin to being tortured, though the child would otherwise have a good life. It is not permissible to conceive in this case, even though the child overall might well have a life worth living. Some things are so awful that one should simply refrain from causing them, and a year of total agony is among those things.

But many people think that being deaf is not so bad. It does not involve pain, and society makes a number of accommodations that make the lives of the deaf easier than they would otherwise be. Thus, to some people, causing deafness may seem to be a kind of harming that can be justified.

On my view, there is a threshold somewhere: some harms are so bad that nothing can justify them, such as the harm of living a year in total agony; other harms are not so bad that nothing can justify them, such as the harms in an ordinary life. My proposal to explain why the hard cases are hard cases is that people implicitly place the threshold between these cases in different places. Some people implicitly see deafness as not so bad, and as a kind of harm that can be justified. Others implicitly see deafness as quite bad and not as a kind of harm that can be justified. If there really is a threshold, it makes sense that there would be implicit disagreement as to exactly where it lies.¹⁰

On a variant of my view, *other reasons* also play a role in justifying procreation. (On my view, they do not.) Supposing that other reasons *do* play a role in justifying procreation, it is nevertheless implausible that *certain kinds* of other reasons can justify harming. Consider a woman with a permanent condition such that if she ever conceives, her child will suffer from a chronic disease that involves some pain but otherwise allows a full life. It is natural to think that there is a lot at stake for this woman: if she does not ever conceive, then she does not get to have a child. This seems to provide a significant

reason in favor of conceiving. However, note that this woman might live in a society like the U.S. where adoption is possible. Thus, what's at stake for the woman is not raising a child, but simply conceiving and gestating her own child. Furthermore, if the permanent condition is genetic, then the woman could have someone else's egg donated and then she could gestate a child without passing on the disease. Supposing that either option is available to her, *what the woman gives up* in failing to conceive (with her own egg) is not *being a parent*, and not even *being a gestational mother*, but simply *being a genetic parent* to her child. I will not deny that there is value in being a genetic parent to one's child; however, I am skeptical that there is sufficient value in this relationship to warrant harming someone to achieve it. And I think similar considerations cast doubt on the significance of being a gestational parent: it seems insufficiently valuable to play a role in justifying harming someone.

A more compelling case involves a woman for whom adoption and egg donation are not available and who has a permanent condition that would cause any child she conceives to have a disease. What is at stake for this woman—being a parent at all—is a big deal. Furthermore, when we think of the question whether slaves in 1800 may permissibly procreate, we certainly imagine the case to be like this: their only chance of having children at all would be to have children that would become slaves. This makes it more likely—but by no means clear—that conceiving is permissible. Add in the further fact that birth control would not be available to them, and they would be giving up even more—that is, having sex at all—if they chose to avoid procreation.

Nevertheless, if any of these other reasons do justify procreation then it's permissible for the parents to harm another person, the created child, in order to avoid

some hardships for themselves. It's not in general permissible to harm others in order to avoid hardships for oneself—when the others would not be in any way guilty or involved in one's experiencing those hardships. So I deny that other reasons ever play a role in *justifying* procreation. (It is much more plausible that other reasons play a role in *excusing* procreation: that they explain why one wrongful act of procreation is less bad than another wrongful act of procreation, or why one agent is less blameworthy than another.¹¹)

4. HARMING AND BENEFITING THOSE WHO INDEPENDENTLY EXIST

I will now turn to discussing some objections to the view I have offered. The first objection is that my view treats knowingly conceiving a deaf child as much worse than it really is. The objector points out that on my view, if one knowingly conceives a child who will be deaf, then one harms the child by causing his deafness. However, the objector points out, what one does is *less bad* than an action which *deafens* an existing child. Yet on my view, both actions harm the child by causing the child to be deaf. It appears that on my view, both actions harm in the same way, and both are thereby equally bad.

This objection brings out something important about the nature of harming in non-identity cases, which is often overlooked. What is brought out is that harming in non-identity cases is in one important respect *less bad* than harming of people who already exist—and in general, harming of people who exist independently of the action. When one harms someone who independently exists, there are two strikes against the action:

-- one causes bad things for the harmed person

-- one prevents the harmed person from an existence without these bad things

These two things are so closely connected in typical cases of harming that they are not separated, but they should be separated. Non-identity cases involve only the first aspect: one harms by causing the bad things for the harmed person, but one does not prevent the person from an existence without these bad things. I think this is a way in which harming someone who independently exists is *worse* than harming someone in a non-identity case: the harm is worse for the independently-existing person, and the action is thereby worse.¹²

These considerations explain some important facts. Many people think that it is wrong to procreate if one has a permanent condition that will result in any child one conceives being deaf. Nevertheless, these people will acknowledge that someone who deafens an existing child is doing something much worse *to the child* than someone who conceives knowing the conceived child will be deaf. On my view, the first action is worse because it harms in two ways: by causing the deafness *and* by depriving the child of being hearing.

Thus, my response to the objector is that while two actions may both harm in the same way, in that they both cause deafness, one action may be worse because it also harms in another way: it deprives the child of being hearing.

The discussion I have offered of harms to existing persons sets me starkly opposed to Derek Parfit's view of the two medical programs case.¹³ In this case, we are considering two different medical programs which would each reduce the number of cases of a certain disability by 1,000. One program would work by pre-conception

screening of women for a temporary condition that would cause the disability in any child they conceived. The other program would work by screening pregnant women for another condition and treating them to prevent the disability from arising in their fetuses. Parfit claims that intuitively it is irrelevant which screening program we use—it is irrelevant that the latter program affects already-existing beings, while the former program affects *who* exists. On my view, that is wrong. If we use the pregnancy-screening program, we are doing more good than if we use the pre-conception screening program. The pregnancy-screening program prevents the disability in independently existing people, while the pre-conception screening prevents the existence of people who would have the disability. Thus, on my view, the pregnancy-screening program prevents worse harm.

The second objection is to claim (4):

- (4) If an action harms someone who does not independently exist, then the fact that the action also benefits the person harmed, and benefits him more than it harms him, is *ineligible* to justify the harm if *failing to perform the action* would similarly benefit someone.

The objector argues that claim (4) is false, as follows. Claim (4) is a restricted claim about some harms only; it is restricted to harms to people who do not independently exist. The objector complains that this restriction is unmotivated and ad hoc: either the general claim is true, in which case the restricted claim is trivially true as well; or the general claim is false, in which case the restricted claim must be false as well. Furthermore, the objector maintains, the general claim is false. There is nothing that could motivate the

restricted claim if the general claim is false. Therefore, the objector concludes, the restricted claim (4) is false as well.

The general claim is this:

- (4*) If an action harms someone, then the fact that the action also benefits the person harmed, and benefits him more than it harms him, is *ineligible* to justify the harm if *failing to perform the action* would similarly benefit someone.

(Note that I do not endorse (4*.) The objector argues that (4*) is false as follows. Suppose that I am in a position to distribute a benefit; it can only go to one person, and I must choose who gets it. Suppose further that Ben and Sally both want it, and that getting it would have no bad effects on Ben, but if Sally gets it she will be harmed—suppose she will suffer a migraine. If the benefit is substantial enough, Sally may want me to choose her despite the fact that getting the benefit would cause her to get a migraine. But if (4*) is true, then I cannot give it to Sally: there is a strong reason against harming her, and the fact that my action would also benefit her is ineligible to justify the harm because I have an alternative that would similarly benefit someone. The objector points out that it clearly would be permissible to give the benefit to Sally (if the benefit is significant enough). The action harms Sally, but the harming is justified *in part* by the fact that it also greatly benefits her. (While the fact that Sally consents to be harmed is also part of the justification of the harm, consent alone is not a sufficient justification for harming someone.)

I think this case does establish that (4*) is false. But it helps to motivate the possibility that (4) is true while (4*) is false. What strikes us about the case is that if we

don't harm Sally, then Sally *loses out* on the benefit. What moves us to think it may be permissible to harm Sally is the thought of what is bad for her about our *not* harming her. This brings out an important difference between actions that affect whether people exist and those that do not. If an action benefits someone who exists independently of the action, then there are two kinds of considerations in favor of the action:

- the action benefits the person
- if the action is not performed, the person is deprived of the benefits (that is, she exists but doesn't get some benefits she could have had)

When a non-identity action benefits someone, only the first of these two considerations is available to tell in favor of the action; the second does not apply. But the second seems to play a crucial role in how, sometimes, the fact that we benefit someone can play a role in making harming that person permissible—and it plays a crucial role in why (4*) is false. Since the second consideration doesn't apply to non-identity actions, this suggests that the reasons (4*) is false just don't apply to (4).

(Note that in discussing the second objection, I haven't offered an independent argument for the truth of (4). Rather, I've defused an objection to it.)

Consideration of these two objections brings out the following. On my view, the fact of *non-identity* in non-identity cases has been widely misunderstood. It is not as significant as some people have claimed: it does not undermine the assertion that harms occur in non-identity cases. But it is not as irrelevant as some have claimed: it does lessen harms (though not very much), and it does limit the significance of benefits: it is harder for benefits in non-identity cases to justify harm.¹⁴

5. IS CAUSING HARM SUFFICIENT FOR HARMING?

In this section, I will discuss another objection to the account I have offered. It is an objection to my claim that one harms someone whenever one *causes* the person to be in a bad state. This objection is an adaptation of an objection to my 2004 and Shiffrin's 1999 from Matthew Hanser's paper in this volume. The objector claims that sometimes one causes someone to be in a particular bad state, but one does not harm the person, because one causes the bad state simply by *improving* a worse state. Improving a bad state, the objector claims, is not a way of harming someone. Here is an example. Suppose that a patient is legally blind; his eyesight is severely impaired, so that he cannot really make out anything visually. Say that his state of severe visual impairment is state S (for "severe"). Suppose that the only thing that can improve this patient's sight is an operation that will improve his sight so that he is somewhat impaired; he will be able to distinguish shapes and colors, but not fine details. This is a significant impairment but much better than his current state. Say that his visual impairment after the operation is his being in state M (for the "milder" impairment). The objector maintains that on my view, a doctor who performs the operation *harms* the patient by *causing* him to be in state M, which is a bad state. However, the doctor does not harm the patient at all; he simply improves his situation. There is no harm-based reason against the doctor's action to be overridden in this case. The objector grants that if there *were* a harm-based reason against acting, it would be overridden, but denies that that tells the right story about the case: this is not a case in which it is permissible to cause one harm because one is preventing (or alleviating) a greater harm. Rather, this is a case in which one simply improves someone's situation.¹⁵

I will say several things in response to this objection. First, it matters what the details are: it matters *how* the operation improves the person's state from S to M. Suppose the operation works as follows. The only way to improve the patient from being in state S is to start a causal chain that would completely cure the sight: using a laser to reshape the retina would create normal sight. However, a side effect of using such a powerful laser would be to distort the nerves that transmit retinal images to the brain. Thus, the causal route by which S is cured is distinct from the causal route by which M is caused. In this case, I claim, the surgeon causes the patient to be in state M; furthermore, the surgeon causes the patient to be visually impaired. In this case, the surgeon does *harm* the patient by causing M, and the case is properly understood as one in which some harm is caused but that is justified because a greater harm is alleviated.

The objector would respond that he imagined the causal structure of the case differently. Suppose the operation works as follows. The patient's sight is initially so poor because his retina has a thick film of fatty tissue over it; this fatty tissue deflects light. The best that can be done is to remove *some* of the fatty tissue. If more were removed, the retina would be damaged. This is how the surgeon improves the patient's state to state M. In this case, M is caused simply as the end-state of improving from state S; M is not caused by an independent process from the process of improving S. Here, the objector's claim has force: it is natural to say that the doctor has not *harmed* the patient at all, but has simply lessened an existing harm. But the doctor does cause the patient to be in state M, a bad state. Thus the objector maintains that, on my view, the doctor harms the patient.

My response to this objection is as follows. In this version of the case, while the doctor has caused the patient to be in state M, the doctor has not caused the patient to be visually impaired. Nor does the doctor cause the patient to be disabled. Now I will clarify the view I stated earlier. In claim (1), I offered a *sufficient* condition for harming. Here is a clarified restatement of that claim:

- (1') One harms someone if one causes him to be in pain, to be in mental discomfort, to be in physical discomfort, to have a disease, to be deformed, to be disabled, or to die.

In this version of the case, because the doctor does not cause the patient to be disabled, this condition is not met. Thus, I am not committed to the view that the doctor harms the patient.

Note that typically, when one causes someone to be in a particular bad state, one also causes him to be in the more general bad state of which this is a particular form. So typically this distinction is not relevant. Furthermore, in non-identity cases, the non-identity actions do cause the created people to be in the more general bad states they are in; non-identity actions do not cause particular bad states without causing the more general bad states.

At this point the objector might revise his objection as follows. Consider the following case. An angry man gets mad at someone in a bar, and strikes him with a broken bottle. Suppose that the victim was already somewhat visually impaired (suppose he was in state L, for “low visual impairment”) but that the angry man’s action increased his visual impairment (moving it to state M). Suppose further that the angry man worsens the victim’s eyesight not by creating a new problem but simply by exacerbating

an existing one. In this case, the angry man causes the victim to be in state M, but does not cause the victim to be visually impaired, and does not cause the victim to be disabled.¹⁶ So my sufficient condition for harm is not met. Yet, the objector might say, not only has the angry man clearly harmed the victim, but he has harmed him in the very way that I am trying to capture with claim (1'): it is the same kind of harm, and it is a failing of my view if my view cannot acknowledge that. The complaint is that the previous version of the doctor case is importantly different, as to whether the agent has harmed, from the angry man case—yet my view does not see a difference between them.

What the objector is bringing out is that there is another sufficient condition for harm besides (1') (and there may be others¹⁷). While the doctor and the angry man both cause people to be in state M, and both do not cause those people to be visually impaired, there are other differences between them. We might offer this sufficient condition:

One harms someone if one *worsens* a bad condition the person was already in.

Some cases of worsening are already covered by (1')—cases in which the way the bad condition is worsened is sufficiently distinct from the original way it was bad—but some, like the case of the angry man, are not.

Perhaps there is a better way to state the new sufficient condition. We might note that, as some have argued,¹⁸ some causal facts are irreducibly of the form: *c* caused *e* rather than *f*. The doctor caused his patient to be in state M rather than in state S. But the angry man caused the victim to be in state M rather than in state L. This is a difference in the causal facts about their actions. We might offer this sufficient condition:

One harms someone if one causes her to be in a particular bad state rather than in a better state.

In summary, it's true that one sometimes causes someone to be in a particular bad state without harming him. My view is that one harms if one causes someone to be in certain general bad states (listed in (1')). Furthermore, there are other sufficient conditions on harming; I have suggested two different ways we might state sufficient conditions to cover the angry man case.¹⁹

6. CONCLUSION

In this paper I have aimed to accomplish eight things.

I have developed a view that elaborates the claim that we have strong reasons against harming; I have offered one way in which we might see the reasons given by harming as *strong* and not easily outweighed.

I have argued for a new taxonomy of non-identity cases:

- (i) same-number cases
- (ii) different-number cases, more people if the action is not performed
- (iii) different-number cases, unknown which scenario has more people
- (iv) different-number cases, more people if the action is performed

I have argued that cases of types (i) to (iii) are very similar in what makes the actions wrong, while cases of type (iv) are quite different.

I have argued that *epistemic* considerations play a crucial role in why cases of type (iii) are wrong.

I have offered an account of why cases of type (iv) are wrong.

I have offered an explanation of why there is disagreement about cases of type (iv).

I have argued that harming independently existing persons is more morally serious than harming in non-identity cases (but not much worse), and also that benefiting independently existing persons has stronger, different reasons in favor of it than benefiting in non-identity cases.

Finally, I have responded to an objection from causing someone to be in a bad state by improving her condition.²⁰

Princeton University

eharman@princeton.edu

REFERENCES

- Hanser, M. *forthcoming* Harm and harming: the non-identity problem revisited, in *Harming Future Persons*, Melinda Roberts and David Wasserman, eds., Springer.
- Harman, E. 2004. Can we harm and benefit in creating? *Philosophical Perspectives* 18: 89-113.
- Schaffer, J. 2005. Contrastive causation. *Philosophical Review* 114 (3): 327-358.
- Shiffrin, S. 1999. Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory* 5 (2): 117-148.
- Parfit, D. 1984. *Reasons and persons*. Oxford: Clarendon Press.

NOTES

¹ A *solution* would vindicate our pre-theoretic beliefs about these cases. Some have argued that no solution is possible, that we can only explain how non-identity actions are wrong by impersonal explanations that do not involve the claim that the actions harm anyone. For example, Parfit 1984 makes this claim.

² I don't claim that everyone has the intuition that these cases are wrong; I simply claim that these intuitions are widespread. And, in particular, they are much more widely shared than are intuitions about what I call the "hard cases" below.

³ The view I offer in this paper is a revision of the view in my 2004. The view is broadly the same, but differs importantly in two ways: my view of the interaction between reasons against harm and reasons to benefit has changed, and my view of what to say about "different number cases" (cases in which one's action affects how many people exist) has changed.

This paper is not meant to supercede but rather to supplement my 2004. In that paper, I take up many questions and argue for many conclusions which I do not address in this paper.

⁴ In my 2004 I argued for (3) by considering two cases that support the claim that actions can be wrong in virtue of harming someone though they also benefit the person more than they harm him. These are not non-identity cases.

Shiffrin 1999 has a view similar to my claims (1) and (3). On her view, reasons to prevent harm are much stronger than reasons to benefit. On her view, preventing someone from worse harm can justify harming her, but providing positive benefits to someone does not similarly justify harming her.

⁵ Note that whether a person who is harmed by an action exists independently of the action is not a matter of what *time* the person exists: if *whether* the person exists does not depend on *whether* the action is performed, the person exists independently of the action, even if the person does not yet exist at the time of the action.

⁶ We need the further condition that the agent knows that *not performing the action* is morally permissible. This condition is satisfied in non-identity cases.

⁷ Note that I'm not assuming that the agents in these cases *know* the moral fact that if this is a different-number case such that more people would exist if the action is not performed, then acting would be wrong. The principle I'm appealing to does not require such knowledge. (Compare: suppose a person knows that if he shoots at random into the forest, there's a 50% chance he'll hit a person, because there are a lot of people wandering around the forest. Then he should not shoot. We needn't add the further claim that he knows shooting with 100% certainty of hitting someone would be wrong.) Someone might object that a theory of what's *really* going on morally in non-identity cases should be an account of what the agents *objectively ought to do*, not of what they *subjectively ought to do*, whereas by appealing to the agents' ignorance, I am really telling a story about what the agents *subjectively ought to do*. The answer to this worry is that our "ought" claims fall along a long scale from *very subjective* to *very objective* and often what we want to explain is what an agent with "ordinary full information" should do in a particular case. But "ordinary full information" rarely contains full information about how all subtle future contingencies will work out, so these "ought" claims are not fully objective. And that is the kind of "ought" claim I am explaining.

⁸ In my 2004, I discussed a type of non-identity different-number case in which more people exist if the harming action is performed. I saw these cases as challenging for my view, but I focused on the idea that the existence of *additional happy people* besides the harmed people put pressure on us to say these actions are permissible, despite the fact

that they harm. I now see that these cases need not have troubled me: the fact that some additional happy people would exist cannot justify harming other people.

⁹ The following objection might be raised. If someone has decided to conceive a baby, and then is considering conceiving *now*, her alternative is conceiving *later*. But then on my view, the benefits cannot justify the harms, because the alternative would involve benefiting someone the same way. My response is as follows. When deciding *whether to conceive*, such considerations do not apply: there is no such alternative. Once one has decided to conceive, conceiving now is no better or worse than conceiving later: either way, one both harms and benefits someone and would have harmed and benefited someone in the alternative. This brings out something important: the fact that both acting and failing to act would harm someone in the same way is a way of justifying that harm.

¹⁰ My explanation of the hard cases involves the following thought: if my view is correct, then someone who had the moral facts *roughly* correct would believe that non-identity actions are wrong in the cases that aren't hard cases, but if she didn't have the moral facts *precisely* correct, then she might get the hard cases wrong. On the assumption that people tend to have the moral facts roughly correct, but don't always get them precisely correct, my view provides an explanation of disagreement in the hard cases: some people are implicitly placing the threshold in the wrong place. That my view can provide an explanation of the fact of disagreement about hard cases counts in favor of my view. I don't claim, however, that this provides anything like a decisive argument for my view. Nor do I claim that in general we're entitled to assume that people are getting the moral facts roughly correct.

There is much discussion nowadays of the question: are we entitled to rely on moral intuitions about specific cases in making moral arguments? My view is that arguments that rely on specific moral claims about cases are not thereby illegitimate as arguments; any argument has undefended assumptions, or starting places, and these may be specific or general. Furthermore, I am unsympathetic to the thought that the existence of disagreement about specific claims undermines arguments that rely on those claims: these arguments have the specific claims themselves as premises, *not* the claim that the specific claims are believed by everyone.

This paper is aimed not to provide a conclusive argument for my overall view, but simply to lay out a view with a certain structure and show that the view has certain advantages. Some of these advantages are: getting the moral facts right. For example, my view implies that it is wrong for a woman to conceive if she has a temporary condition that will cause her child to be blind; and it implies that her action is wrong in virtue of harming her child. When I say that it's an advantage of my view that it gets the moral facts right, I'm relying on claims like this: what the woman does is wrong. I don't mean to be supporting that claim by the further claim that everyone agrees with that. But another part of my argument for my view—and it is only a small part of the support for my view—is the claim that my view can explain why there is disagreement about some cases while there is agreement about other cases. It is here, and only here, that I intend my view to be supported by claims about what people believe.

¹¹ See my 2004.

¹² Someone might object that what I claim is true of harms to independently-existing people *in general* is only true of some harms. Consider a case of preemption: Joe punches Steve in the face, but Bill would have swung at Steve if Joe had not. The objector claims that while Joe harms Steve by *causing* him pain and injury, Joe does not *prevent* Steve from an existence without this pain and injury, because if Joe had not acted, Steve would still have suffered in the same way. My response is that the objector's claim is false: Joe does prevent Steve from an existence without this pain and injury. Steve *is prevented* from an existence without this pain and injury. Who prevents him? Joe does. (Bill, who stands idly by, certainly does not.) So, we see that *prevention* does not require counterfactual dependence: one can prevent an outcome even though the outcome would not have occurred had one not acted.

¹³ Parfit (1984), pp. 366-369.

¹⁴ Shiffrin 1999 points out that if one benefits someone by creating her, then if one had not acted, she "[would] not experience [the benefit's] absence" and it's not the case that her life will "go worse" (134). But Shiffrin does not think a parallel consideration is relevant to harming: "Harms may differ [from benefits]—a harm that does not make a life worse does not always have less or significantly less weight than harms that do worsen lives, marking another asymmetry between harms and benefits" (footnote 32). By contrast, I think that the fact that harming an independently existing person deprives her of life without that harm does make such harming, at least to some degree, worse than harming by creating.

One might object to my account as follows. I claim that it's particularly hard for benefiting by creating to justify harm in non-identity cases, because benefiting by creating is less morally significant than other benefiting. The objector points out that on my view, harming by creating is *also* less significant than other harming. Perhaps a less significant harming can be justified by a less significant benefiting. My response is that my goal in this paper is to spell out a certain kind of view—a view on which we take seriously that there are strong reasons against harming. Thus, on the view I am spelling out, there are strong reasons against all kinds of harming, including harming by creating, even though this is somewhat less bad than harming independently-existing persons.

¹⁵ Note that what's at issue between the objector and me is not whether the ordinary word "harm" applies but, taking harm to be distinctively serious and to generate strong harm-based reasons, whether the doctor has harmed the patient. In fact, I will readily admit that on my view some actions harm though we would not typically say they harm; my claim is that these actions generate strong harm-based reasons. (For one example, if I affect when a couple conceives by keeping them late at dinner, my action is a cause of the existence of the particular child they create, thus my action harms the child (and also benefits the child). But while on my view, such actions do harm, there is no practical upshot of the fact that they harm. When we engage in everyday actions like this, some of which affect whether people come to exist, we don't know whether performing the action is more likely to harm than not performing it, and so harm-based reasons aren't available to us in our reasoning and don't play a role in what we should choose to do.)

¹⁶ It's not clear to me whether the causal facts can ever be this way: whether one can *worsen* someone's eyesight without it's also being true that one has caused the person to

be visually impaired. But I will grant this to the objector. If it's not possible, then my claim (1') handles this case directly.

¹⁷ For example, some would say that I harm someone if I deprive her of something to which she is entitled (or has a right). I am not taking a stand on whether our reasons against depriving someone of a positive good—a benefit—to which she is entitled are *strong* in the way that harm-based reasons against causing pain and other bad things are strong.

¹⁸ Schaffer (2005).

¹⁹ In his paper in this volume, Hanser also makes the following objection to my 2004. Suppose that two people are in danger and I can save exactly one of their lives; however, to save one of them I would have to break his arm. Hanser points out that the fact that I would benefit him (by saving his life) *is* sufficient to justify harming him. He takes this to be an objection to my claim that the mere fact that one greatly benefits someone, and benefits him more than one harms him, is not sufficient justification for harming him. However, this is not a case of both harming and benefiting someone, as I use the terms “harm” and “benefit.” What I mean by “benefit” is to provide a positive good to someone, such as a good experience. This is a case in which it is permissible to harm someone because one is *preventing* greater harm. (My view on this matter is the same as the view Hanser attributes to Shiffrin.)

²⁰ For helpful feedback on drafts of this paper, I am grateful to Tyler Doggett, Peter Graham, Alex Guerrero, Sarah McGrath, Brad Skow and audiences at University of Vermont and Ohio State University.