

Chapter 9

Harming and Procreating

Matthew Hanser

Abstract Suppose that A harms B if and only if he stands in some appropriate relation to B's suffering of some harm. Such an account of harming can explain how it is that those who choose Derek Parfit's risky energy policy harm the future people later exposed to radiation, but can it also yield the result that parents who conceive a disabled child, when they could easily have conceived a different, non-disabled child instead, thereby harm the child they actually conceive?

Keywords Harm · Non-identity problem · Wrongful life · Derek Parfit · Seana Shiffrin · Elizabeth Harman.

9.1

Derek Parfit argues that a choice does not harm someone in a morally relevant sense if that person has a life worth living and would not have existed had the choice not been made.¹ This creates a puzzle regarding choices that affect the identities of future people. Consider two examples:²

The Risky Policy. We must choose between two long-term energy policies, one of which involves the burial of radioactive waste in an area that we know will become earthquake prone in the distant future. We choose this policy, and three centuries later an earthquake releases radiation that injures or kills thousands of people. Because this policy also raises the standard of living, the people injured or killed by the leaked radiation would never have existed had we chosen a safer policy. (Standards of living have trivial but cumulative effects on people's daily lives, eventually affecting when and with whom people have sex. A choice that affects the standard of living thus affects the identities of future people.)

The Disabled Child. A couple knows that because the man carries a genetic defect any child they conceive in the ordinary way will suffer from a fairly serious disability. They also know that their doctor can isolate a sperm from the man, genetically alter it to eliminate the defect, combine it *in vitro* with one of the woman's eggs, and implant the resulting zygote in the woman's uterus. Although they could easily afford this procedure (perhaps it is covered by

M. Hanser (✉)
University of California, Santa Barbara, CA, USA
e-mail: hanser@philosophy.ucsb.edu

their insurance), they choose to conceive a child unaided. Had they opted for the *in vitro* procedure they would almost certainly have had a different child, since the sperm and egg used in that procedure would almost certainly have been distinct from the ones that gave rise to their actual child.

If Parfit is right, our choice of the risky policy does not harm the people injured or killed by radiation, since they have lives that are on balance worth living and would not have existed had we chosen a safer policy. Likewise, the couple's choice to conceive unaided does not harm their child, since he too has a life worth living and would not have existed had they selected the *in vitro* procedure. Why, then, are these choices morally objectionable? Parfit calls this the *non-identity problem*, since it arises because the people who live short, unhealthy, or disabled lives as a result of these choices are distinct from the people who would have lived longer, healthier, or disability-free lives had the choices not been made. Parfit argues that we cannot give a *narrow person-affecting* account of what's objectionable about these choices—we cannot give an account that appeals only to the choices' (supposed) bad effects on people who actually ever live. In each case the objection must be either (i) that the choice's effects on the people it brings into existence are not as good as the effects that some other choice would have had on the people *it* brought into existence, or (ii) that the choice produces a worse overall state of affairs, impersonally conceived, than some other choice would have produced.

In an earlier paper (1990) I argued that a choice *can* be objectionable owing to its bad effects on someone who has a life worth living and who would not have existed had the choice not been made. I argued further that on a correct analysis of what it is for a choice to harm someone, the choice of the risky policy *does* harm the people injured or killed by leaked radiation. But I was reluctant to extend this analysis to cases like that of the disabled child.³

More recently, Seana Shiffrin (1999) and Elizabeth Harman (2004) have presented accounts of harming structurally quite similar to my own. But they take their accounts to apply quite straightforwardly to procreative choices: neither would hesitate to say that the parents of the disabled child harm him by bringing him into existence. So who is right here? Was my reluctance to extend my account of harming to cases like that of the disabled child misplaced, or is there a genuine problem here, the force of which Shiffrin and Harman have failed to appreciate? That is the question I wish to explore in this paper.

Here's another way of framing the issue. I've described two cases, that of the risky policy and that of the disabled child. Parfit thinks that in each case there is an obstacle (the same obstacle) to saying that the choice harms someone. Shiffrin and Harman think that in neither case is there an obstacle to saying this. I think there is no obstacle in the case of the risky policy but that there may well be one in the case of the disabled child. (Obviously I must disagree with Parfit about what this obstacle is—it can't simply be that the child has a life worth living and would not otherwise have existed.) For the purposes of this paper I shall assume that I am right and Parfit wrong about the risky policy; I don't want to reargue that here. The question addressed in this paper is whether there is some obstacle to saying that the parents harm their child *other* than the one (supposedly) identified by Parfit.

9.2

Let's begin with Parfit's obstacle and why I think it is illusory. The obstacle is created by the claim that an agent does not harm someone by acting in a certain way if that person has a life worth living and would not have existed had the agent acted differently. (Parfit prefers to speak of a choice's harming someone rather than of an agent's harming someone by acting in a certain way, but I do not think that anything hangs on this.) Parfit defends this claim by appealing to a certain general account of harming. Put roughly, the account is this:

A harms B by acting in a certain way if and only if B would have been better off had A acted in some other way.

For present purposes, the "only if" direction is the one that matters: A harms B by acting in a certain way *only if* B would have been better off had A acted in some other way. This is why Parfit thinks we do not harm the people injured or killed in the nuclear catastrophe. These people have lives worth living and would not have existed at all had we chosen some other policy. They consequently would not have been better off had we chosen some other policy.

There are a number of problems with this account of harming. For one thing, it should probably be revised to say that A harms B by acting in a certain way if and only if B would have been better off *in some respect* had A acted in some other way. An agent can harm someone in a morally relevant sense even though he also benefits him, and even though the benefit outweighs the harm. In such a case the agent's action is still *pro tanto* objectionable; it may even be objectionable all things considered. (*Pro tanto* objectionable acts are ones that stand in need of justification; in the absence of justification they are objectionable all things considered.) But this revision does not suffice to undermine Parfit's claim about the risky policy, for there is arguably *no* respect in which the people exposed to radiation would have fared better had we chosen some other policy. In order for there to be a respect in which someone is better off in one scenario than in another, he must have some level or other of well-being in each scenario. And in order for a person to have a level of well-being in a scenario, he must exist in that scenario. (A person who does not exist does not have a "neutral" level of well-being.)

The real problem with Parfit's account of harming is deeper than this. The account is defective because it attempts to analyze the notion of harm all in one go, so to speak. We should instead adopt a two-step account. The first step is to analyze the notion of *harming* (i.e. of *doing harm*) in terms of the prior notion of someone's *suffering harm*; the second step is then to say what it is for someone to suffer harm. Schematically:

Step One: A harms B by acting in a certain way if and only if A's acting in that way results in his standing in relation R to some harm suffered by B.

Step Two: B suffers harm if and only if ____.

To transform this account-schema into a substantive account of harming we must do two things: specify relation R, and fill in the blank in the account of what it is to

suffer harm.⁴ In my earlier paper I did not offer a substantive account of harming. I gave relation R a name—I said that A harms B if and only if his actions make him at least partially *responsible* for a harm suffered by B—but I did not attempt to spell out the conditions for an agent's being at least partially responsible for a harm.⁵ Nor did I offer an account of what it is for someone to suffer harm. Rather, I argued that on *any* reasonable account of what it is to suffer harm, the people exposed to radiation thereby suffer harm, and that on *any* reasonable account of what it is for an agent to be responsible for someone's suffering harm, we who adopt the risky policy are responsible for the harms these people suffer. The harms they suffer, after all, are the predictable result of the release of radiation; and the release of radiation is the predictable result of our burying radioactive waste in an area which we know will eventually become earthquake-prone. On any reasonable way of fleshing out the two-step account of harming, then, we who adopt the risky policy harm the people later exposed to radiation.

9.3

Since our choice of the risky policy predictably harms people, it is *pro tanto* objectionable. But it does not follow that the choice is objectionable all things considered. To secure this stronger result we need to show that the choice is not an instance of *justified* harming. And a case can be made that it is. Despite the harms they suffer, life is on balance good for the people exposed to radiation. Since they would not have enjoyed this good but for our choice of the risky policy, we benefit them by choosing this policy. And since the benefit we bestow upon them is greater than the harm we cause them to suffer, our choice is not objectionable all things considered. Or so the reasoning goes.

Shiffrin, Harman and I all reject this line of reasoning, but our arguments are different, and it is worth pausing briefly to examine the differences. In my earlier paper I raised two objections.⁶ (In that paper I did not, and in this section I do not, question the assumption that being conceived can come to one as a benefit.) First, I expressed doubt about the assumption that if a choice brings someone more benefit than harm, then the harm-based *pro tanto* objection to the choice is automatically overridden. But instead of developing and defending this doubt, I concentrated on a second objection. I argued that by choosing the risky policy we do not, in a morally relevant sense, benefit the people who would not otherwise have existed. My claim was really that we cannot use the benefits they receive to *justify* our choice. We cannot do this because only *intended* benefits can be used to justify harming; and I argued that when we choose the risky policy we cannot plausibly be taken to intend the benefits received by those who will later owe their existence to our choice. But while I think it is defensible that only intended benefits can be used to justify acts of harming, this is a claim that many would reject. Dialectically speaking, then, I was perhaps unwise to base my argument upon it.

Harman's argument is simpler.⁷ She argues, in effect, that since by stipulation we will benefit an equal number of future people whichever policy we choose,

the benefit-based reasons in favor of choosing the competing policies cancel out, leaving the harm-based *pro tanto* objection to choosing the risky policy intact. This argument is unsatisfying. First, it can be deployed only with respect to same number choices—choices where the same number of people will ever live whichever choice is made. If a choice will affect the number of future people, then the benefit-based reasons favoring the competing options do not completely cancel.⁸ More importantly, the argument does not accurately reflect the way justification works. Suppose that I can save one or another of two people from death, but not both. And suppose that while I can save the first person without harming him, I can save the second only at the cost of breaking his arm. Clearly it would be permissible for me to save the second person. The benefit I can bestow upon him only by breaking his arm *justifies* me in breaking his arm, which is to say that it justifies me in harming him. This justification is not undermined by the fact that I could have bestowed an equivalent benefit upon someone else without doing harm. If Harman's reasoning were sound, saving the second person by breaking his arm would be impermissible: the benefit-based reasons in favor of saving each of the two people would cancel out, leaving the harm-based *pro tanto* objection to harming the second person intact.

Shiffrin has perhaps the most persuasive argument why the harm-based *pro tanto* objection to choosing the risky policy is not overridden.⁹ She distinguishes two kinds of benefit. Someone receives what we might call a *preventative* benefit when he is prevented from suffering harm. (For Shiffrin, to suffer harm is to be in a certain sort of noncomparatively bad state. I shall say more about this in the next section.) Someone receives a *pure* benefit, by contrast, when he receives a positive good—a good that does not consist in the absence or prevention of something bad. When someone is prevented from falling down a manhole, he receives a preventative benefit; when someone who is comfortably well off wins the lottery, he receives a pure benefit. Shiffrin argues that it is generally permissible to harm someone, even without his consent, in order to prevent him from suffering a greater harm. It is permissible, that is, to harm someone, even without his consent, in order to bestow upon him a more-than-compensating preventative benefit. That is why it is permissible to break a person's arm (even without his consent) when that is the only way to save his life. But Shiffrin argues that it is not generally permissible to harm someone, without his consent, merely to bestow upon him a more-than-compensating pure benefit. We cannot, for example, break someone's arm, without his consent, in order to help him win the lottery. Now if someone with a life worth living receives a benefit by being brought into existence, that benefit is pure, not preventative. (Coming into existence does not prevent one from suffering the "harm" of nonexistence.) We consequently cannot justify harming someone by pointing out that he would never have existed had we acted differently.¹⁰ (As a number of authors have argued, that the person would not later regret our action is neither here nor there.¹¹)

Shiffrin's argument applies both to same number choices and to different number choices. It also applies both to our choice of the risky policy and to the parents' choice to create a disabled child. Suppose her argument is sound. Then *if* the parents harm their disabled child by bringing him into existence, they cannot claim that the

harm-based *pro tanto* objection to their choice is overridden by the fact that he would not otherwise have received the benefit of existence. Consequently, if there is an asymmetry between our choice of the risky policy and the parents' choice to create a disabled child, it must consist in this: that while we harm people by choosing the risky policy, the parents do *not* harm their disabled child by bringing him into existence.

By contrast, if my argument is sound and Shiffrin's is not, there are potentially two asymmetries between these choices: perhaps the parents do not harm their child by bringing him into existence; and if they do, perhaps that harming is justified. It would be justified if existence were a benefit that the parents *intended* to bestow upon their child.¹² But I shall set the question of justification aside for the remainder of this paper. I shall concentrate solely on the question whether the parents of the disabled child harm him by bringing him into existence.

9.4

Recall that according to the two-step account of harming, A harms B by acting in a certain way if and only if A's acting in that way places him in relation R to some harm suffered by B. If there is an obstacle to saying that the parents harm their disabled child, then, it is either because there is an obstacle to saying that their child suffers harm or because there is an obstacle to saying that the parents' action places them in relation R to whatever harm he suffers. Now I claimed in Section 9.2 that on any reasonable account of what it is to suffer harm, the people exposed to leaked radiation suffer it, and that on any reasonable account of relation R, we who choose the risky policy stand in relation R to the harms they suffer. That is why Shiffrin, Harman and I can agree that the choice of the risky policy harms people without having first to agree upon how exactly the two-step account is to be fleshed out. But things are different in the case of the disabled child. Whether the child suffers harm depends upon which particular account of harm we adopt; and whether the parents stand in relation R to that harm, assuming that it exists, depends upon how exactly we understand relation R. In this section I shall consider whether the disabled child suffers harm. In subsequent sections I shall address the question whether the parents' actions connect them in the right way to that harm, assuming that it exists.

Of course like everyone else, the disabled child can be expected to suffer numerous harms over the course of his life. He'll suffer harm, for example, when he's beaten up by a playground bully in the fourth grade. I shall consider later whether parents, simply in virtue bringing their children into existence, thereby become at least partially responsible for all the harms their children ever suffer. But the example of the disabled child is meant raise a narrower question: do the parents harm their child by bringing him into existence *with a disability*? The question to be addressed in the present section is thus whether the disabled child suffers harm simply in having, or in coming into existence with, a disability.

The most widely accepted account of harm (call it the *standard* account) relies upon a counterfactual comparison:

A person suffers harm if and only if there occurs an event such that he would have been better off in some respect had it not occurred.

This account of harm must be distinguished from Parfit's superficially similar account of *harming*. According to Parfit, A harms B by acting in a certain way if and only if B would have been better off (in some respect) had A acted differently. The risky policy can be used to illustrate the difference between the two accounts. According to Parfit's account of harming, we do not harm anyone by choosing the risky policy, because the people exposed to radiation would not have been better off (in any respect) had we acted differently. But according to the standard account of harm, the people exposed to radiation do suffer harm: they would have lived longer, healthier lives had the event of their exposure to radiation not occurred.¹³

So let's apply the standard account of harm to the case of the disabled child. First we must distinguish between, on the one hand, the *actions* that may or may not place the parents in relation R to some harm suffered by their child, and, on the other hand, the *event* (whatever it might be) that may or may not come to their child as a harm. Perhaps this event is the child's conception. Of course we sometimes speak as if a child's conception were an action, but strictly speaking it's a "mere" event, consisting (let us suppose) in the joining together of two gametes. This event may be caused by actions performed by the parents, but it is not itself an action. Does the event of the disabled child's conception come to him as a harm? According to the standard account, it does not. Had the child not been conceived, he would not have lived a life free of disability—he would not have been in a better state with respect to the ability in question. Rather, he would not have existed at all. And we may imagine the details of the example to be such that there is *no* event such that had it not occurred, that very same child would have been born free of disability. According to the standard account, then, the child suffers no harm with respect to his disability.

But what if we accept some other account of harm? Shiffrin, who argues vigorously against comparative models of harm, proposes that harms be identified with "certain absolute, noncomparative conditions (e.g. a list of evils like broken limbs, disabilities, episodes of pain, significant losses, death) . . ."¹⁴ To suffer harm, according to Shiffrin, is not to be in a comparatively bad condition; it is to be in a certain sort of *absolutely* bad condition. It is not to be worse off than one was before, or than one otherwise would have been; it is simply to be *badly* off in a certain way. And she offers an account of what unifies the various items on her list: harms are "conditions that generate a significant chasm or conflict between one's will and one's experience, one's life more broadly understood, or one's circumstances."¹⁵ The details of her account, however, needn't concern us. The important point is that on her account, the child's disability is itself a harm. That he has always had this disability and that he could not have existed without it are irrelevant to the question whether he suffers harm in having it.

Harman agrees that disabilities are harms even for those who have always had them and who could not have existed without them. According to Harman, states of ill health, bodily damage and bodily deformity count as harms because it is worse to be in one of these states than to be in a healthy, normal state. The states on her list are comparatively bad, but the states to which they compare unfavorably are those of a healthy, properly functioning organism of the relevant species. It doesn't matter whether the individual suffering harm was ever in a position to attain this ideal. On Harman's account as well, then, the disabled child suffers harm simply in being disabled.¹⁶

I do not accept the standard account of harm, but nor do I accept Shiffrin's or Harman's account. On my view, harms paradigmatically consist in losses of "basic goods." The power of sight, for example, is a basic good for a human, so *losing* the power of sight is a harm. Merely *lacking* the power of sight, by contrast, is not a harm.¹⁷ If this is right, then someone who comes into existence lacking the power of sight has not thereby suffered harm: one can't lose what one has never had. Likewise with other disabilities. So here there is indeed a difference between my view and the views of Shiffrin and Harman. They think the disabled child suffers harm simply in being disabled, and this leaves room for the possibility that his parents harm him by bringing him into existence. On my view, the disabled child does not suffer harm simply in being disabled.¹⁸ Nor does he suffer harm in coming into existence with a disability. There is thus no room to say that the parents harm him by bringing him into existence with a disability.

But this difference in our views is more verbal than real.¹⁹ While I do not think the child suffers harm in being disabled, I do think his disability is *bad* for him—his disabled state is one that it is bad for him to be in. And I think that if his parents' actions place them in relation R to his disability, then their actions are *pro tanto* objectionable.²⁰ Shiffrin, Harman and I thus agree that if the parents' actions relate them in the right way to their child's disabled state, their actions are *pro tanto* objectionable. We differ only in whether we would in that case describe the parents as having *harmed* their child. For the sake of argument, then, I shall for the remainder of this paper accept Shiffrin's and Harman's description. I shall grant that disabilities are harms.

9.5

Shiffrin and Harman hold that a correct understanding of harming will reveal that the parents' actions satisfy the very same sufficient condition for harming that paradigm ordinary acts of act of harming satisfy. (By "ordinary" acts of harming I mean ones where the victim would have existed even if the action had not been performed.) The parents' actions may differ from ordinary acts of harming in certain respects, but the differences are irrelevant: it's what they all have in common that makes them harmings. That, at any rate, is the idea. Now according to the two-step analysis of harming, A harms B by acting in a certain way if and only if A's acting in that way places him in relation R to some harm suffered by B. I have granted, for the sake

of argument, that the disabled child suffers harm in being disabled. So finally we arrive at the crux of the matter. Do the parents' actions place them in relation R to their child's disability?

In order to answer this question, we need to know what conditions are sufficient for the obtaining of relation R. Shiffrin tells us that her argument begins with "the (contested) assumption . . . that being created can harm a person," and she explains in a footnote that the basic idea behind this assumption is that "by being the direct cause of a person's being in a situation that intrinsically delivers harm, one harms that person. . . ." ²¹ Likewise, Harman writes that "[a]n action harms a person if the action causes pain, early death, bodily damage, or deformity to her . . ." ²² Both, then, suggest that *causation* of harm is sufficient for harming. And both treat it as obvious that the parents of the disabled child cause his disability.

According to what is probably the most popular account, causation is the ancestral of the *counterfactual dependence* relation, with the relata being distinct, actually occurring events: event *e* depends counterfactually upon event *c* if and only if *e* would not have occurred had *c* not occurred; and *c* *causes* *e* if and only if the two events are connected by a chain of such counterfactual dependencies. ²³ (The chain might of course consist of a single link.) But this understanding of causation will not suit Shiffrin's and Harman's purposes, since they are interested in the causation of states, not the causation of events. (For both Shiffrin and Harman, harms are bad states or conditions.) Let us assume that a state is a thing's possession of a property. We may adapt the counterfactual account of causation accordingly:

β's being *G* depends counterfactually upon *α*'s being *F* if and only if *β* would not have been *G* had *α* not been *F*;

α's being *F* causes *β*'s being *G* if and only if the two states are linked by a chain of such counterfactual dependencies.

This account yields the desired result that the parents' bringing their child into existence causes him to be disabled: had they not brought him into existence, he would not have been disabled.

The parents harm their child in the morally relevant sense if (a) this account of causation is correct and (b) causation of harm is sufficient to generate a harm-based *pro tanto* moral objection. ²⁴ We could attempt to evaluate these two claims separately, but this would involve us in difficult, disputed questions about the nature of causation. ²⁵ I thus propose collapsing the two claims and evaluating the result directly. The collapsed claim is this:

If B's suffering some harm is linked, via a chain of counterfactual dependencies, to A's acting in a certain way, then A's acting in that way is *pro tanto* objectionable.

I shall argue that this claim is false. Perhaps this means that causation is not the ancestral of counterfactual dependence; perhaps it means that causation of harm is not sufficient to generate a *pro tanto* moral objection; perhaps it means both. We needn't decide.

The harm we've been discussing so far has been that of a child's being disabled. But being disabled is a *determinable* property. There are many dimensions along

which someone can be disabled, each corresponding to an aspect of proper functioning; and along each dimension there are a spectrum of ability levels, ranging from fully normal functioning (or even better than normal functioning) to a complete lack of functioning. Even a single power, such as sight, can be assessed along multiple dimensions. Let us choose a particular dimension of functioning for our example: the ability to discriminate detail visually at a distance. At one end of the ability range is vision that is 20/20 or better. Moving away from this ideal we find increasing degrees of shortsightedness—vision that is 20/30, 20/40, and so on. Having precisely 20/200 vision is a fully *determinate* disability.²⁶

Now suppose a doctor performs an operation to improve his shortsighted patient's vision. Before the operation the patient's vision is 20/250; afterwards it is 20/200. According to Shiffrin's and Harman's accounts of harm, someone with 20/200 vision suffers harm; so even after the operation, the patient suffers harm. But the doctor obviously does not harm his patient by performing the operation. There is no harm-based *pro tanto* objection to the doctor's action. In order for an action to be *pro tanto* objectionable, then, it is *not* enough that had the action not been performed, some person would not have been in the particular state of harm he's actually in. Nor is it enough that the person would instead have been in another state from along the relevant ability spectrum.²⁷ What matters is whether he would have been in a *better* state from along that spectrum had the action not been performed. If the doctor were to make his patient's sight *worse*, he would harm him. But since he actually *improves* his patient's sight along the relevant dimension, he benefits rather than harms him.²⁸

Shiffrin and Harman could, I suppose, insist that the doctor *does* harm his patient by performing the operation, and then argue that this harming is *justified*. The doctor causes his patient to suffer the harm of having precisely 20/200 vision, but he also prevents him from suffering (or continuing to suffer) the worse harm of having precisely 20/250 vision. The doctor thus inflicts a lesser harm upon his patient in order to prevent him from suffering a greater harm—a paradigm case of justified harming! But this is not a plausible analysis of the example. To be sure, there is such a thing as inflicting a lesser harm for the sake of preventing a greater one. Suppose that the only way a doctor can cure his patient of a disease causing permanent paralysis is to give him a drug that will, as a side effect, make him shortsighted. Here the doctor clearly harms his patient by giving him the drug. He damages his patient's sight, which we may suppose would otherwise have been perfect. But this harming is justified. The doctor inflicts harm along one dimension of functioning in order to prevent greater harm along another dimension. In the operation example, however, only one dimension of functioning is involved. The doctor simply decreases his patient's degree of shortsightedness. That he was unable to eliminate the shortsightedness completely—that he was unable to improve his patient's vision all the way to 20/20—is not a “cost” or unfortunate side effect of the operation, to be weighed against the benefit. In this example the doctor *simply* benefits his patient; he does not harm him *at all*.

In denying that the doctor harms his patient, I am not rejecting noncomparative accounts of harm. I agree that the patient's state after the operation is still one that it

is bad for him to be in. For the sake of argument I have even granted that in having 20/200 vision the patient suffers harm. What I am denying is that the doctor's action *relates* him to this harm in a way that makes his action *pro tanto* objectionable. Nor am I abandoning the two-step analysis of harming. I am not returning to Parfit's position that an agent harms someone only if he makes him worse off than he would otherwise have been. To determine whether the doctor harms his patient we do not compare levels of well-being; we compare degrees of *harm* along a single dimension of functioning. It is significant, however, that it has turned out to be impossible to avoid comparisons in giving an account of harming. If we eliminate them from our account of what it is to suffer harm, they reemerge in our account of what it is to stand in relation R to harm.

The claim we wished to evaluate was:

If B's suffering some harm is linked, via a chain of counterfactual dependencies, to A's acting in a certain way, then A's acting in that way is *pro tanto* objectionable.

We have seen that this is false. In order to determine whether an agent stands in relation R to a person's having some determinate disability, it is not enough to know that the person would not have had that particular, determinate disability had the agent acted differently. We must also know whether he would have been in a better or a worse state of functioning along the relevant dimension had the agent acted differently. The comparison is crucial: if the person would have been in an even worse state of functioning had the agent acted differently, the agent benefits rather than harms him.

If the claim I've rejected has seemed true, I suspect this is because people writing on the issue have tended to focus upon certain *determinable* disabilities. While determinate disabilities occupy single points along ability spectra, determinable disabilities disjunctively encompass whole ranges of ability levels. Someone is *mildly shortsighted*, for example, just in case his determinate level of visual acuity falls below 20/20 but above (let us suppose) 20/60. Since this range includes neither end of the relevant ability spectrum, the fact that a mildly shortsighted person would not have been mildly shortsighted had a certain agent acted differently is compatible both with its being the case that his vision would have been better had the agent acted differently and with its being the case that his vision would have been worse had the agent acted differently; and of course if the subject's vision would have been worse had the agent acted differently, the agent benefits rather than harms him. But things are different when the range of ability levels disjunctively encompassed by the determinable disability includes the low end of a spectrum. One such disability is being legally blind. Someone is legally blind just in case his determinate level of visual acuity falls below a certain threshold. If someone who is legally blind would not have been legally blind had a certain agent acted differently, this *rules out* the possibility that his vision would have been even worse had the agent acted differently. It consequently rules out the possibility that the agent benefits rather than harms him with respect to that dimension of functioning.²⁹ Even more extreme is the determinable harm of being disabled *simpliciter*. Someone is disabled *simpliciter*

just in case there is some dimension or other along which his level of functioning falls below the threshold of proper functioning. As long as someone exists, the only alternative to his being disabled *simpliciter* is his functioning properly in every respect.

This does not mean that how alternative levels of functioning compare is irrelevant to the question whether an agent stands in relation R to someone's being legally blind, or to his being disabled *simpliciter*. It just means that so long as the victim was going to exist either way, the outcome of the comparison is not in doubt: if the subject would still have existed but would not have been legally blind (or disabled) had the agent acted differently, then the subject would have functioned *better* along the relevant dimension had the agent acted differently; and so the agent harms him by acting as he does. There is no reason to think that whether B would have functioned better had A acted differently has any less bearing on whether A stands in relation R to B's being legally blind, or disabled *simpliciter*, than it does on whether A stands in relation R to B's being mildly shortsighted, or having precisely 20/200 vision. It would be absurd to suggest that a comparison between levels of functioning is relevant to harming only when the harm in question happens not to encompass the low end of an ability spectrum.

9.6

I have argued that the mere fact that B would not have suffered some harm had A acted differently is not enough to make A's act one of harming. It matters what the alternative was for B. Now in the procreation example the disabled child would not have existed at all had his parents not created him. Might the fact that *this* was the alternative for their child be enough to make their action one of harming?

The argument of the previous section does not rule this out. We've seen that if B, who has a certain determinate disability, would have been in an even worse state with respect to that ability had A acted differently, then A does not harm him, but rather benefits him by acting as he does. And we may presumably add that if B would have been in exactly the same state of disability had A acted differently, then (setting aside worries about preemption) A neither harms nor benefits him. So perhaps A harms B as long as neither of these conditions obtains—perhaps A harms B if B is in a disabled state and would have been in neither that same state nor a worse one (along the same dimension) had A acted differently. This sufficient condition for harming could in turn be satisfied in either of two ways: by its being the case that B would have been in a better state (along the relevant dimension) had A acted differently or by its being the case that B would not have existed at all had A acted differently.

The resulting view would be this. Suppose that B has some determinate disability. Then:

1. A *benefits* (and does not harm) B with respect to the relevant dimension of functioning if B's state of functioning along that dimension would have been even worse had A acted differently;

2. A *neither benefits nor harms* B with respect to the relevant dimension of functioning if B's state of functioning along that dimension would have been neither better nor worse had A acted differently;
3. A *harms* B with respect to the relevant dimension of functioning if *either* (i) B's state of functioning along that dimension would have been better had A acted differently, *or* (ii) B would not have existed at all had A acted differently.

I don't think this view is acceptable. The aim of the approach to the non-identity problem under discussion in this paper is to show that non-identity cases *aren't special at all*, that the acts of the parents and of those who choose the risky policy satisfy the *very same* sufficient condition for harming that paradigm ordinary acts of harming satisfy. The putative sufficient condition considered and rejected in the previous section met this *desideratum*. The condition stated in (3), by contrast, is disjunctive: it has separate clauses for identity and non-identity cases. And the non-identity disjunct—clause (ii)—seems completely ad hoc. All the other conditions specified in the view—the condition for benefiting, the condition for neither benefiting nor harming and the condition for harming in identity cases—rest upon *comparisons* between actual and counterfactual states of functioning. Why should there be a second, *noncomparative* condition for harming? This condition appears to have been tacked on simply to give the desired result in non-identity cases.

There is a further respect in which clause (ii) seems ad hoc. Consider a sequence of cases: first, an agent causes someone to have 20/200 vision rather than 20/190 vision; next, he causes someone to have 20/200 vision rather than 20/180 vision; and so on, the alternative each time approaching closer to 20/20. As the sequence continues, the acts of harming become progressively more serious. Eventually we reach the most serious case of all, the one in which the victim would have had perfect eyesight had the agent acted differently.³⁰ Now according to the view characterized by (1)–(3), equally serious (presumably) would be the harming done if the person with 20/200 vision would not have existed at all had the agent acted differently. But why should we assimilate nonexistence to perfect vision when determining whether the agent harms the person who ends up with 20/200 vision? Why not assimilate nonexistence to a complete lack of sight? Surely we could with equal justice replace (1) with:

- 1'. A *benefits* (and does not harm) B with respect to the relevant dimension of functioning if *either* (i) B's state of functioning along that dimension would have been even worse had A acted differently, *or* (ii) B's actual state of functioning along that dimension is at least better than a complete lack of functioning and B would not have existed at all had the agent acted differently.

My opponent says of the child with 20/200 vision: "Had his parents not created him, he would not have suffered the harm of being shortsighted. It's not just that he wouldn't have suffered the determinate harm of having 20/200 vision. He would not have suffered harm of *any* degree along the relevant dimension. There is consequently a respect in which his parents harm him by creating him." But we can with equal justice say: "Although it is bad for the child that he has 20/200 vision, at least he can see to *some* degree; had his parents not created him, he would not have been

able to see to *any* degree. His parents consequently benefit him by creating him. It's not just that they benefit him on balance, since his life is good on balance. They benefit him even with respect to his power of sight."³¹

Now in fact I don't think we should say either of these things. Nonexistence should be assimilated to neither end of the ability spectrum. Comparisons are necessary for determinations of harming, and while 20/200 vision can be compared with other degrees of visual acuity, it cannot meaningfully be compared with nonexistence.³² I suspect my opponent's point, that the disabled child would not have suffered harm of any degree along the relevant dimension had his parents acted differently, gets its rhetorical force from cases where comparisons *are* possible—from cases where the victim would have existed even if the agent had acted differently. In such cases the victim would have been in a good state, and hence in a *better* state, along the relevant dimension had the agent acted differently.

Alternatively, perhaps the reason some have thought it sufficient for harming that the person who suffers harm would not have existed had the agent acted differently is that they have thought it sufficient for harming that the person in question would not have suffered that particular, determinate harm had the agent acted differently; and satisfaction of the former condition entails satisfaction of the latter one. But as we saw in the previous section, the latter condition is not sufficient for harming. We require *independent* grounds for thinking that the former condition is sufficient for harming, and so far we have found none.

9.7

In Section 9.5 we considered a condition satisfied both by the parents of the disabled child and by the agents of most ordinary acts of harming: had they acted differently, certain people would not have suffered the harms they actually suffer. Unfortunately, satisfaction of this condition turned out to be insufficient for harming. In Section 9.6 we considered a condition satisfied both by the parents of the disabled child and by the agents in other non-identity cases, but not by the agents of ordinary acts of harming: had they acted differently, certain people who suffer harms would never have existed. But the suggestion that satisfaction of this condition suffices for harming turned out to be unmotivated. Now let's consider a condition the satisfaction of which *does* suffice for harming. The question will be whether it is satisfied by the parents of the disabled child.

Our reflections on the eye operation example suggest the following sufficient condition for harming:

A harms B by acting in a certain way if A's acting in that way rather than in some other way helps explain why B is in a certain bad state (with respect to a given dimension of functioning) *rather than in a better state* (with respect to that same dimension of functioning).³³

This condition appeals to what we might call "contrastive" explanation. For present purposes, we may understand contrastive explanation as the ancestral of

contrastive counterfactual dependence: β 's being G rather than G' depends counterfactually upon α 's being F rather than F' if and only if α is F rather than F' , β is G rather than G' , and if α had been F' rather than F , β would have been G' rather than G .³⁴

Now it is not true that if the parents of the disabled child had acted differently, their actual child would have been in a better state of functioning. They consequently do not satisfy this sufficient condition for harming via a *single-linked* chain of counterfactual dependencies. But we must consider whether their performing some action rather than another helps explain, via a *multi-linked* chain of counterfactual dependencies, why their child is in a certain bad state (with respect to a given dimension of functioning) rather than a better one. This possibility must be taken seriously, for it is only in virtue of such a multi-linked chain that we who choose the risky policy harm those who die in the ensuing nuclear catastrophe. After all, it is not true that if we had chosen some other policy, the people actually killed by leaked radiation would have lived to ripe old ages. There is no single-linked explanation here either. Working backwards, the explanatory chain runs as follows: the people's dying prematurely rather than living to ripe old ages depends counterfactually upon the earthquake zone's having radioactive waste buried in it rather than not having radioactive waste buried in it; and the earthquake zone's having radioactive waste buried in it rather than not having radioactive waste buried in it depends counterfactually upon our having chosen the risky policy rather than some other policy.

Can we construct such a multi-linked chain in the case of the disabled child? Let us try, again working backwards. The child would have functioned properly had the sperm from which he partially originated not been genetically defective. (We will suppose, perhaps implausibly but for the sake of argument, that this genetic defect was not identity determining.) The child's being disabled rather than fully functional, then, depends counterfactually upon the sperm's being defective rather than non-defective. And the sperm's being defective rather than non-defective in turn depends counterfactually upon the father's being a carrier of the defect rather than a non-carrier. No doubt we could continue tracing this chain of counterfactual dependencies back in time. The problem is that no *action* of the father's need appear in the chain. There need be nothing the father did, for example, such that the sperm's being defective rather than non-defective depends counterfactually upon his having done that rather than something else. Likewise, there need be nothing the father did such that his being a carrier of the genetic defect rather than a non-carrier depends counterfactually upon his having done that rather than something else. And so on.

We can of course imagine a version of the example in which the father's having acted in one way rather than another *does* figure into the relevant explanatory chain. His being a carrier of the defect rather than a non-carrier, for example, might depend counterfactually upon his having taken a certain drug in his youth rather than abstaining. In that case he *would* have harmed the child he later helped create. But there is no reason why the child's being disabled rather than fully functional *must* be explained by either parent's having at some point acted in one way rather than another.³⁵

9.8

Let us assume, then, that no action of either parent figures into the chain of counterfactual dependencies explaining why their child is disabled rather than fully functional. Mightn't their act of creating him figure into the explanation in some other way? Had they not created him, this explanatory chain would never have been forged, for in that case there would have been no child whose condition needed explaining. By creating the child, then, the couple gave the fact that the man was a carrier of a genetic defect rather than a non-carrier the *opportunity*, so to speak, to explain someone's being disabled rather than fully functional. (I continue to assume, for the sake of argument, that the genetic defect is not identity-determining.) The parents knew that if they created a child, the man's being a carrier of the defect rather than a non-carrier would lead to that child's being disabled rather than fully functional. Mightn't that be enough to make their act one of harming?³⁶

I don't think so. Suppose that B has 20/200 vision. A arranges for C to perform an operation on B; C is the only doctor in the world capable of performing this sort of operation. Unfortunately, C is also a misogynist, and although he would never admit it, he never does his best work on women. C performs the operation on B, improving her vision to 20/150. Had he not been a misogynist, he would have improved it to 20/100. B's having 20/150 rather than 20/100 vision after the operation is thus explained by C's being a misogynist rather than a non-misogynist. And by arranging for C to perform the operation, A gave C's misogyny the opportunity, so to speak, to explain B's having 20/150 rather than 20/100 vision. But A's action is not even *pro tanto* objectionable. A does not harm B by arranging for C to perform the operation. (C arguably harms B by not performing the operation as well as he should have done, but this feature of the example is not essential. We could imagine instead that the operation's effectiveness was reduced by some temporary environmental condition. Unfortunately, the operation could not be delayed.)

In this example, of course, B would have been even more shortsighted had A not arranged for the operation, whereas in the procreation example the child would not have existed at all had his parents not created him. But we have yet to see a reason for thinking that this difference *makes* a difference. It would be ad hoc to suggest that while there is no *pro tanto* objection to giving a condition or circumstance the opportunity to explain someone's being disabled rather than fully functional when that person would otherwise have been in an even worse state of functioning, there *is* a *pro tanto* objection when either (i) the person would otherwise have been in a better state of functioning or (ii) the person would otherwise not have existed at all.

9.9

Here's another possibility. That his parents created him rather than doing something else does not explain why the child is disabled rather than fully functional, but it does explain why he is existent rather than nonexistent, alive rather than not alive.

His parents thus stand in relation R to his existence or his life. But perhaps, in virtue of standing in relation R to his existence or his life, they indirectly also stand in relation R to each of the elements (states, episodes, etc.) making up that existence or life. Some of these elements are good, others are bad; and among the bad elements is his disability. Because his parents *create* him—give him life—they stand in relation R to everything in his life, including his disability. And so even though his life is on balance good, they harm him by creating him.

I think this proposal's plausibility rests upon equivocations. When we speak of life, we sometimes mean the sort of thing that a person can live or lead: a long life or a short one, a happy life or a sad one, an easy life or a hard one, a boring life or an exciting one. In this sense, a person's life is something like his history. But we can also mean the bare property of being alive. Likewise, when we speak of existence, we sometimes mean the sort of thing that can be pleasant, carefree, meager or wretched, but at other times we mean the bare property of existing. Now the parents' creating their child rather than doing something else does *not* explain why their child has one sort of life or existence rather than another. There is thus no reason to think that the parents stand in relation R to their child's life or existence in the first sense. (It won't do to say that the child would not have had the sort of life or existence he actually has if he had not been created—we've already seen that that's not sufficient for the holding of relation R.) All we're entitled to say is that the parents stand in relation R to their child's having the bare property of being alive, or of existing: their creating him rather than doing something else explains why he is alive rather than not alive, existent rather than nonexistent. The bare property of being alive, or of existing, however, does not have both good and bad "elements." I thus see no argument from the parents' standing in relation R to their child's being alive, or existing, to their standing in relation R to the various good and bad elements that comprise their child's life or existence.

9.10

I agree that being disabled is bad for the child, and that this fact must figure somehow into a correct account of why his parents' choice to create him is *pro tanto* objectionable. One way it could figure in is this: perhaps the parents produce a worse outcome by creating a disabled child (or one who suffers harm) than they would have produced by creating a child free of disability (or one who didn't suffer harm). But for a variety of reasons I would prefer a less impersonal explanation of why the parents' choice is *pro tanto* objectionable. Of course if the parents were to *harm* their child by creating him, that would provide a satisfyingly person-affecting explanation, but for reasons I have attempted to set out in this paper I have serious doubts about the availability of this explanation. Even if we broaden the concept of *harming* for moral purposes, as I think we should, the concept must still pick out a relation between agents and harms that is intuitively sufficient to generate a harm-based *pro tanto* objection, and I have yet to find convincing grounds for

thinking that such a relation holds between the parents of the disabled child and their child's disability.

Acknowledgements I am grateful to David Wasserman and Melinda Roberts for their helpful comments.

Notes

1. Parfit (1984), pp. 372–374. Parfit distinguishes between “morally relevant” and “ordinary” senses of harming. Using the word in its ordinary sense, it might be infelicitous to say that an agent harms someone if his action is only a remote cause of that person's suffering harm, but Parfit points out that such an action might nonetheless be morally objectionable owing to its role in the production of that harm. I agree that for purposes of moral theory we should employ a somewhat broader notion of harming. Henceforth I shall usually drop the modifying phrase “in a morally relevant sense,” taking it as understood. Parfit has a second reason for distinguishing morally relevant from ordinary harming: an action that harms someone in the ordinary sense needn't be morally objectionable. I find this reason unconvincing. An act of *justified* ordinary harming is not morally objectionable all things considered, but it is still, in a morally relevant sense, an act of harming. That's why it stands in need of justification. The relevant question is whether ordinary harmings ever require no justification at all; and I do not find Parfit's examples convincing on this score. See Parfit (1984), pp. 69–71.
2. The first example is a simplified version of one found in Parfit (1984), pp. 371–372. The second example comes from McMahan (1998), p. 208.
3. Nor did I defend what is sometimes called the person-affecting *restriction*: roughly, that an action cannot be wrong unless there is someone whom it harms, wrongs, or causes to be worse off.
4. When I say that Shiffrin, Harman and I offer structurally similar accounts of harming, I mean that our accounts share this two-step structure. We all think that A harms B if and only if he's connected in the right way to a harm suffered by B.
5. By “responsible for harm” I did not mean “at fault for, or to blame for, harm.” I meant that the harm is attributable to one's agency in such a way that one can legitimately be asked to defend one's role in bringing it about. If one's conduct makes one at least partially responsible for a harm, one's conduct stands in need of justification. The important question is what non-moral relation between one's action and a harm can make one at least partially responsible for that harm.
6. Hanser (1990), pp. 60–62.
7. Harman (2004), pp. 92–93.
8. Harman grants that if the risky policy would result in a larger future population, the benefit-based reason in favor of adopting it would outweigh the benefit-based reason in favor of adopting the competing policy. But she argues that the moral reason against harming is so serious that adopting the risky policy would still be impermissible. See Harman (2004), pp. 102–103.
9. Shiffrin (1999), pp. 119–133. Her actual argument concerns procreative choices, but it applies equally well to the choice of the risky policy.
10. Strictly speaking, Shiffrin does not argue that it is *impermissible* to harm someone (without his consent) in order to bestow upon him a greater pure benefit. Her article concerns the question whether the agents of such acts *owe compensation* to those they have harmed; her conclusion is that they do. By contrast, an agent who unavoidably harms someone in the course of saving his life has no obligation to compensate the recipient of his beneficence.

Shiffrin also argues that the force of the reason in favor of bestowing a pure benefit is

greatly reduced when no one will be left in a comparatively bad state if one fails to bestow that benefit. If a child is not conceived, his life does not go worse than it would have gone had he received the benefit of existence. The fact that a certain choice is a necessary condition for someone's receiving the benefit of existence consequently has limited justificatory force. See Shiffrin (1999), pp. 134–135.

11. See, for example, Woodward (1986), pp. 822–825, and Harman (2004), pp. 98–101.
12. This is close to being the view of David Wasserman (2005).
13. According to the standard method for evaluating counterfactuals, a given counterfactual is true if and only if its consequent is true in the “closest” possible world—the world most similar to the actual world—in which its antecedent is true. The world most similar to our own in which the people are not exposed to radiation is presumably one in which the earthquake does not occur, but whose history leading up to that point is otherwise the same. It is *not* a world in which we chose some other energy policy three centuries earlier. David Heyd accuses me of equivocating when I say both that the people who die in the nuclear disaster would have been better off had they not been exposed to radiation and that they would not have been better off had we had chosen a different policy. Heyd (1992), p. 113. But when I say this I am not, as he supposes, illicitly describing a single event first as the people's exposure to radiation and then as our choice of the risky policy. Rather, I am describing two different events which occur three hundred years apart.
14. Shiffrin (1999), p. 123.
15. *Ibid.*
16. Harman (2004), pp. 96–97.
17. I defend this understanding of harm, and criticize competing conceptions, in Hanser (2008). I should stress that my view is not a version of the temporal comparison view. I do not hold that suffering harm is a matter of being worse off in some respect than one was at an earlier time.
18. I remain open to the possibility that the child's disability is a *harmful* condition—a condition that gives rise to harms. But it is not itself a harm.
19. I do not mean that nothing of significance hangs on which account of harm is correct. I mean only that my disagreement with Shiffrin and Harman over whether disabilities are harms does not affect the question whether we can give a narrow person-affecting account of why it's wrong (assuming that it is wrong) to create a disabled child, when one could easily create a non-disabled child instead.
20. In my earlier paper I argued that if an action makes its agent responsible for someone's being in a bad state, then even if the victim cannot be said to suffer harm as a result of the action, the action is *pro tanto* objectionable on narrow person-affecting grounds. See Hanser (1990), pp. 64–65.
21. Shiffrin (1999), p. 119 and fn. 8. She continues that “it is sufficient for this article, though, to assume that if one is the direct cause of a person's being in a situation that intrinsically delivers harm, then one is responsible for that person's suffering harm, even if one does not *harm* her.” I take this last phrase to mean “even if one does not harm her *in the ordinary sense.*”
22. Harman (2004), p. 93. And regarding a case like that of the disabled child, she writes that “conceiving harms the child because it causes the child to be in a particular kind of bad state.” *Ibid.*, p. 94.
23. This account derives from Lewis (1973).
24. Some might object that causation of harm generates a *pro tanto* objection only when the harm is foreseeable. Since the parents could foresee that their child would be disabled, I shall set this worry aside.
25. For an overview of the pros and cons of various counterfactual analyses of causation, see the papers collected in Collins et al. (2004).
26. Given the possibility of corrective lenses, shortsightedness, unless extremely severe, is not a serious disability, and I'm not sure to what extent Shiffrin and Harman would consider it a harm. But this is not important. I use shortsightedness as an example only because there is a familiar scale for measuring its degrees.

27. This might, however, be sufficient for causation. Had the doctor not performed the operation, the patient's vision would have been 20/250 instead of 20/200. Perhaps it follows that the doctor (or the operation he performed) *caused* the patient to have precisely 20/200 vision. But again, for our purposes this doesn't matter.
28. In this argument I have assumed, along with Harman, that states of less than proper functioning are harms. For Shiffrin, harms are (roughly) states that create a cleavage between a person's experience and his will. The argument could easily be recast in these terms: an action that decreases the degree of cleavage (along some dimension) between a person's experience and his will, without eliminating that cleavage entirely, leaves him in a state of harm but does not harm him.
29. The same point holds for any determinate disability occupying the low end of an ability spectrum—being completely blind, for example. If someone who is completely blind would not have been completely blind had a certain agent acted differently, then it can't be that his vision would have been even worse had the agent acted differently.
30. For convenience I ignore the possibility of vision that is better than 20/20.
31. Would this latter line of thought be as plausible if the harm in question were an episode of pain rather than a disability? One might think not, reasoning as follows. Nonexistence is more plausibly assimilated to the *absence* of something than to its presence; that is why nonexistence is more plausibly assimilated to the complete absence of an ability (i.e. to the *worst* state with respect to the ability) than to its complete presence. But pain is bad in itself—it is not merely the absence of something good. It is thus more plausible to assimilate nonexistence to the complete absence of pain (i.e. to the *best* state with respect to pain) than to a state of maximum pain. I am unconvinced. Episodes of pain are episodes of unpleasant *consciousness*; the end-points of the relevant spectrum are consequently states of pain-free consciousness and states of maximally painful consciousness. Nonexistence can be plausibly assimilated to neither end-point.
32. It is possible to get around this difficulty when the subject's *overall* quality of life is at issue. Suppose a certain sort of life is so bad that it would be better to die now than spend the rest of one's days living such a life. Strictly speaking the terms of comparison here are alternative possible lives—it's better to have a shorter such life than a longer one—but speaking loosely we might say that it's better not to exist at all than to live such a life. Unfortunately, this maneuver cannot be used to defend the claim that the parents of the disabled child harm him by creating him. It simply isn't true that it would be better for the child if he were to die now rather than continue living with 20/200 vision. It isn't even true that *at least so far as his ability to discriminate detail visually is concerned*, it would be better for him if he were to die now rather than continue living with 20/200 vision (if indeed this thought even makes sense). Nor do I think it would be better for a person, at least so far as *pain* is concerned, if he were to die now rather than spend the rest of his life experiencing mild pain (if indeed this thought even makes sense).
33. Since I don't think that harms are bad states, I do not myself accept this as a (partial) account of harming. But I do think that it captures a sufficient condition for an action's being *pro tanto* objectionable owing to its role in explaining someone's being in a bad state.
34. I don't really think that explanation should be cashed out in terms of counterfactual dependence, for the usual reason: in cases of "preemptive" explanation, α 's being F rather than F' explains β 's being G rather than G' even though β would have been G anyway, owing to some other cause, had α been F' . But if we set cases of preemption aside, counterfactual dependence will serve well enough. (I shall not here attempt to determine whether a thing's being F rather than F' amounts to anything more than its being F and not F' .)
35. Suppose a doctor wakens a patient from an otherwise permanent coma, knowing that the patient will suffer mild pain for the remainder of his days. Does the doctor harm the patient? There are two versions of the case to consider. If the pain results from an injury caused by the waking procedure, then I think the doctor does harm the patient. The patient's being in a painful rather than a pain-free conscious state depends counterfactually upon his suffering rather than not suffering the injury; and his suffering rather than not suffering the injury

depends counterfactually upon the doctor's performing rather than not performing the procedure. (This is, of course, a case of justified harming.) But if the pain results from a condition that the patient had even before going into a coma, I am inclined to say that the doctor does not harm the patient. There is nothing the doctor does such that his doing that rather than something else explains the patient's being in a painful rather than a pain-free conscious state upon waking from his coma.

36. Perhaps Harman is suggesting something along these lines when she writes that "[i]f some factor for which we are not responsible will cause a bad result if we act in a certain way, this fact can provide a reason against the action." Harman (2004), p. 95, emphasis omitted.

References

- Collins, J., Hall, N., and Paul, L.A., eds. 2004. *Causation and counterfactuals*. Cambridge, Massachusetts: The MIT Press.
- Hanser, M. 1990. Harming future people. *Philosophy & Public Affairs* 19: 47–70.
- Hanser, M. 2008. The metaphysics of harm. *Philosophy and Phenomenological Research* 77: 421–450.
- Harman, E. 2004. Can we harm and benefit in creating? *Philosophical Perspectives* 18: 89–113.
- Heyd, D. 1992. *Genethics*. Berkeley and Los Angeles: University of California Press.
- Lewis, D. 1973. Causation. *Journal of Philosophy* 70: 556–567. Reprinted in his *Philosophical papers, volume 2*. Oxford: Oxford University Press.
- McMahan, J. 1998. Wrongful life: Paradoxes in the morality of causing people to exist. In *Rational commitment and social justice*, eds. J. Coleman and C. Morris, 208–247. Cambridge: Cambridge University Press.
- Parfit, D. 1984. *Reasons and persons*. Oxford: Clarendon Press.
- Shiffrin, S. 1999. Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory* 5: 117–148.
- Wasserman, D. 2005. The nonidentity problem, disability, and the role morality of prospective parents. *Ethics* 116: 132–152.
- Woodward, J. 1986. The non-identity problem. *Ethics* 96: 804–831.