

## CAN WE HARM AND BENEFIT IN CREATING?

Elizabeth Harman

### 1. The non-identity problem

The non-identity problem concerns actions that affect who exists in the future. If such an action is performed, certain people will exist in the future who would not otherwise have existed: they are *not identical* to any of the people who would have existed if the action had not been performed. Some of these actions seem to be wrong, and they seem to be wrong in virtue of harming the very future individuals whose existence is dependent on their having been performed. The problem arises when it is argued that the actions do not harm these people—because the actions do not make them worse off than they would otherwise be.<sup>1</sup> Consider:

Radioactive Waste Policy: We are trying to decide whether to adopt a permissive radioactive waste policy. This policy would be less inconvenient to us than our existing practices. If we enact the newly-proposed policy, then we will cause there to be radioactive pollution that will cause illness and suffering. However, the policy will have such significant effects on public policy and industry functioning, that different people will exist in the future depending on whether we enact the policy.

Two things should be emphasized. First, the illness and suffering caused will be very serious: deformed babies, children with burns from acid rain, and adults dying young from cancer. Second, the policy will affect who will exist in the future because our present practices invade people's everyday lives, for example by affecting recycling practices in the home; these practices will change if the policy is adopted. Furthermore, whether we adopt the policy will determine which plants are built where, what jobs are available, and what trucks are on the road. These effects will create small differences in everyone's lives which ultimately affect who meets whom and who conceives with whom, or at least when people conceive. This affects who exists in the future.<sup>2</sup>

In the Radioactive Waste Policy case, it seems clear that:

- (1) It is wrong to adopt the Policy because of the harm to the future individuals who will suffer.

And so it must be true that:

- (2) There are reasons against the Policy in virtue of the harm to the future individuals.

However, the following claim may also seem plausible:

Worse-Off: An action harms a person only if it makes the person worse off than she would otherwise have been if the action had not been performed.

And if Worse-Off is true, then (2) is false. The affected people are not made worse off than they would otherwise have been if the policy had not been adopted. If the policy had not been adopted, they would not have been better off, because they would not have existed.

The non-identity problem is the problem of explaining why the actions in Radioactive Waste Policy and similar cases are impermissible. It seems that the actions are wrong, and that they are wrong because they cause harm. However, it is arguable that there is no harm, because Worse-Off can seem plausible.

Some people draw the following lesson from these cases:

Impersonal Explanation: The correct explanation of the impermissibility of the action is not that it harms: it does not harm. Rather, the action is impermissible because the world is *better* if the action is not performed; it is *impersonally* better, though it is not better for any person.<sup>3</sup>

In non-identity cases, it intuitively appears that there is harm and that the harm is morally explanatory of the wrongness of the actions. Those who endorse the Impersonal Explanation consider themselves driven to that explanation by the impossibility of making a harm-involving explanation succeed for these cases. My central goal is to argue that this lesson is not warranted. There is harm in non-identity cases,<sup>4</sup> and this harm explains the wrongness of the actions.

This paper has four goals:

- Goal #1: Solve the non-identity problem by relying on reasons against harm.
- Goal #2: Vindicate the presence and explanatory value of reasons against harm in the cases that generate the non-identity problem.

Goal #3: Argue for a view about reasons to benefit persons.

Goal #4: Argue for a good method of finding out whether there are any reasons against an action due to harm, and argue that an alternative method is misguided.

Achieving the first goal would be valuable in its own right, because people have found it genuinely difficult to explain what is going on in non-identity cases. Achieving the second goal would demonstrate that the Impersonal Explanation is false.<sup>5</sup> The third goal naturally connects to the first two, because it turns out that reasons to benefit persons play an important part of the story about non-identity cases. The fourth goal settles a worry some people have had about non-identity cases, about how we can actually have reasons that stem from harm that would befall people who don't exist.<sup>6</sup>

## 2. My solution to the non-identity problem

The most compelling argument in favor of Worse-Off is based on the following case:

Surgery: A doctor cuts a hole in my abdomen in order to remove my swollen appendix. Cutting open my abdomen causes me pain (as I recover); but if the operation had not been performed, I would have suffered worse pain and died very soon.

In Surgery, the operation is permissible: it is not wrong. Furthermore, it may seem that the operation does *not* harm me: rather, it helps me. Consideration of Surgery naturally leads us to embrace Worse-Off. We appear to need an account of why what the surgeon does to me is not harm; Worse-Off seems to give the right account.<sup>7</sup>

While it might appear to be intuitively obvious that the doctor does not harm me, what is really clear is that what the doctor does is *permissible*. Saying the doctor does not harm me is one way of establishing the permissibility of his action; but it is not the only way. Suppose that instead we say, as I think we should, that the doctor does harm me. He harms me because he causes significant damage to my body. Then there is a reason against performing the surgery in virtue of the harm to me. But consider what reasons there are against *not performing* the surgery: there is a reason against not performing the surgery in virtue of the fact that I would suffer more severely and die if the surgery is not performed. In this case, we can say that the reasons against each course of action are of the right type to be weighed against each other.<sup>8</sup> The reasons against *not performing* the surgery are stronger than the reasons against performing the surgery; and so performing the surgery is permissible.<sup>9,10</sup>

This discussion of the Surgery case suggests a way of handling the Radioactive Waste Policy case.

Just as the surgeon causes significant bodily damage to the patient, and for this reason harms the patient, adopting the Policy causes significant bodily damage to future individuals. I claim that *causing* pain, early death, bodily damage, and deformation is harming. We do not need a complete analysis of what it is to harm, in order to reach this conclusion; we can hold that these are clear cases of harm. We can then say that, for each of these individuals, the fact that adopting the Policy would harm her is a reason against adopting the Policy. So:

- (2) There are reasons against the Policy in virtue of the harm to future individuals.

Now we need an argument from (2) to:

- (1) It is wrong to adopt the Policy because of the harm to the future individuals who will suffer.

It might seem to go as follows: There are compelling reasons against adopting the Policy in virtue of the harm to future individuals. The reason in favor of adopting the Policy, that otherwise we would be inconvenienced, is not sufficient to outweigh these compelling reasons. Therefore, adopting the policy is impermissible.

However, one might object to this argument. Just as adopting the policy causes harm to the future-existing individuals, it also seems to cause benefits to these individuals: every benefit in their lives would not have been had if the policy had not been adopted, because then they would not have existed. This objection holds that the policy benefits the individuals as well as harming them, and that the benefits outweigh the harms, and render the action permissible.<sup>11</sup>

I will argue in section 6 that the objector is right that:

- (3) There are reasons in favor of a course of action in virtue of the benefits to the future individuals who will not exist if the course of action is not taken.

However, we can still argue from (2) to (1). We need to look at all the reasons we have:

Reasons in favor of adopting the Policy:

- if the Policy is adopted, those future individuals who will thereby exist will enjoy many benefits throughout their lives
- if the Policy is not adopted, we will continue to be somewhat inconvenienced

Reasons in favor of *not adopting* the Policy:

- if the Policy is not adopted, those future individuals who will thereby exist will enjoy many benefits throughout their lives
- if the Policy is adopted, those future individuals who will thereby exist will be harmed

The reasons to do with benefit are parallel: they are roughly equivalent. But the reason in favor of not adopting the policy, due to harm, is much stronger than the reason in favor of adopting the policy, due to the inconvenience to us. Here again, the reasons in favor of not adopting the policy are stronger than the reasons for adopting the policy: the balance of reasons tells against adopting the policy.

The central point is the following. Adopting the policy would harm future individuals. Even if we grant that it would also benefit them, adopting the policy is nevertheless wrong because we have an alternative that would involve parallel benefits, without parallel harm.

An objector might point out that there are two types of cases in which the balance of moral reasons tells against a course of action: those in which the course of action is wrong, and those in which refraining from it is supererogatory. I have not shown that adopting the policy would be wrong, merely that the balance of moral reasons tells against it. A first response to this objection is that because my solution succeeds in showing the balance of moral reasons to favor *not adopting the policy* because of the presence of reasons against harm, it already does what proponents of the Impersonal Explanation think is impossible. They think we cannot support the claim that there is harm in these cases and that reasons against harm support an overall moral evaluation that tells in favor of not adopting the policy. Second, the objection can be answered. Giving an account of what makes the difference between requirement and supererogation is beyond the scope of this paper. But in offering my solution, I am making a claim about this difference: reasons against harm are serious enough that the mere presence of “compensating” reasons to benefit those harmed is *insufficient* to render the harms permissible, and is insufficient to make refraining from adopting the policy merely supererogatory in the Radioactive Waste Policy case.

I have now presented the core of my solution to the Non-Identity Problem. My solution involves two crucial claims. First:

- (4) An action harms a person if the action causes pain, early death, bodily damage, or deformity to her, even if she would not have existed if the action had not been performed.<sup>12</sup>

Second, reasons against harm are so morally serious that the mere presence of greater benefits to those harmed is not in itself sufficient to render the harms permissible: when there is an alternative in which parallel benefits can be provided without parallel harms, the harming action is wrong.

In section 9, I will augment my solution to the Non-Identity Problem in addressing a variant of the case in which we affect not only *who* exists but *how many people* exist. In such cases, the benefits that would be bestowed by each possible course of action are not parallel. My solution will rely on taking reasons against harm to be very serious and not easily or straightforwardly outweighed by reasons to benefit those who are harmed. The intervening sections, particularly section 7, will develop this view of reasons against harm.

### 3. Other cases that raise the non-identity problem

My solution applies directly to other cases that raise the Non-Identity Problem. Consider:

**Teenage Mother:** A 14-year-old girl decides to conceive now and raise the child. Because she is so young, she gives her child a bad start in life: her child suffers from inadequate parental stability and support. If she had not conceived now, she would have waited and had a different child later, to whom she would have given a good start in life.<sup>13</sup>

**Temporary Condition:** Tammy has a temporary condition that will cause any baby she conceives now to be born deaf. She conceives now. If she had not conceived now, her condition would have cleared up in two months and she would have conceived a different baby who would not have been disabled.<sup>14</sup>

In both cases, it seems that the woman acts wrongly by conceiving. The Teenage Mother case is disputable on this point; so let's assume that the teenager knows that her child will suffer substantially as a result of her inadequate parenting, in the form of trauma and deep emotional problems.<sup>15</sup> The women do not make their children worse off by conceiving them. The children are not worse off than they would otherwise have been if they had not been conceived: they would not have existed. The problem, again, is to explain why there are reasons against the actions in virtue of the effects on the children, though the actions do not make the children worse off.<sup>16</sup>

In both cases, conceiving causes the resulting child to be in a bad state. I claim that, in each case, conceiving harms the child because it causes the child to be in a particular kind of bad state. In making this claim about the Teenage Mother case, I am expanding on what I have said so far about harm. Suffering the trauma and emotional problems resulting from inadequate parental support, even if it involves no physical injury, constitutes suffering a harm. Thus, the solution I advocate for the Radioactive Waste Policy case can be directly applied

to these cases. Because each woman has an alternative in which she provides parallel benefits without parallel harms, her harming action is wrong.<sup>17</sup>

#### **4. Previous harm-based solutions**

Other writers have also argued that the Non-Identity Problem can be solved by denying what I call the Worse-Off claim and asserting that the future individuals are harmed by the Policy. I will briefly explain why my solution is preferable to the accounts of Hanser (1990) and Woodward (1986). The central problem with both accounts is that they cannot handle the Teenage Mother and Temporary Condition cases.

Hanser (1990) says about a case similar to the Radioactive Waste Policy case that the affected individuals suffer harms as a result of the presence of the nuclear waste, and that because we are responsible for the presence of the waste, we are responsible for these harms. Hanser's view cannot explain the reason against Tammy's action in Temporary Condition. Hanser says, about a case similar to Temporary Condition, that the mother is "not responsible for the defect-producing circumstances of her pregnancy" (69); thus he would say that Tammy is not responsible for the badness to the child of deafness, because Tammy need not be responsible for the condition that will cause her child to be deaf. Hanser sees "wrongful procreation" (69) as raising a special problem; he does not attempt a solution. But surely this is a mistake. If 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. Hanser faces the same problem in handling Teenage Mother: the teenager is not responsible for the fact that she is not older, more mature, and more able to be a good mother.

Woodward (1986) says about a case similar to Radioactive Waste Policy that the affected people "have rights that others should not knowingly pursue policies that will kill, injure, or poison them...at least when there is no weighty justification for such policies and alternative policies which involve no such risks are available" (812). His explanation appears to be complicated by the desideratum of showing, through the violation of the right alone, not just that there is a reason against the Policy, but that this reason is not outweighed; so Woodward builds a consideration about the existence of alternatives into the right he posits. I think this is unnecessarily complex. More importantly, Woodward's account of what rights exist appears to be ad hoc. Woodward (1987, a defense of his 1986) says that surgery, if consented to, does not violate any rights. Because he only considers reasons stemming from rights, he denies my claim that there is a reason against surgery to be outweighed. He makes a number of claims about the presence and absence of rights that do get the correct permissibility results, but he does not do any further explanatory work to explain why the rights exist or do not exist in the particular cases (see 808). For example, he says that there is no right not to

have one's arm amputated if this is necessary to save one's life. But it's not clear why persons would lack that right but have a right not to be poisoned by nuclear waste, even though they could not exist without this harm.

Woodward (1986, 815–816) offers a very different account from my own of why the mother's action in *Teenage Mother* is impermissible. He says that we should not undertake actions that will lead us to have moral duties that we cannot meet. If the teenager knows she will be unable to meet her duties to be a good mother to her child, then she should not conceive. But while this is a part of the explanation of why the teenager's action is impermissible, it misses the most crucial part of the explanation: the harm to the child.

Because Woodward aims to explain the impermissibility of the actions in *Non-Identity Problem* cases through the notion of *rights violations*, he faces a particular problem in one version of the *Temporary Condition* case. Suppose the condition causes deafness by affecting the genes of the egg the mother's ovary releases. While not having one's needs met by one's mother might constitute a rights violation, Woodward does not see any rights violation in this version of *Temporary Condition*, and he cannot account for the wrongness of creation in this case.

## 5. Objection: harm as comparative

One might object to my account by presenting a second argument for *Worse-Off*:

*Worse-Off*: An action harms a person only if it makes the person worse off than she would otherwise have been if the action had not been performed.

Suppose the objector grants that to harm someone is to do something that is bad for her. However, the objector asserts that the notions of something's being good or bad for someone are derivative from the notions of something's being better or worse for someone. The objector maintains that there is no non-relative goodness or badness for someone. So, an action's being bad for a person must consist in its being worse for her that the action be performed than that it not be performed. And so, *Worse-Off* must be true.

I will not here address the difficult question of whether the notions of something's being good or bad for someone are derivative from the notions of something's being better or worse for someone. Even if they are, *Worse-Off* can still be false. Suppose the objector is right that an action is bad for a person just in case the action makes the person worse off than some point of comparison. The objector is mistaken to assume that the only available point of comparison is what things would have been like if an action had not been performed. I propose that for persons, there is a point of comparison that involves a healthy bodily state.<sup>18,19</sup> At least, an action harms someone if it causes the person to be in a state, or to endure an event, that is worse than life with a healthy bodily



state. A healthy bodily state involves no *damage*: no cuts or burns or diseases. But, I claim, it also involves no *deformity*: it is the normal healthy state of an organism of the species in question. So, even a baby that is genetically determined to be deformed is harmed by its deformity.<sup>20</sup> Furthermore, life with a healthy bodily state involves a normal human lifespan; so it does not involve *early death*. I am leaving it vague what counts as a healthy bodily state, but that is okay. The important point is that the kinds of pain, early death, bodily damage, and deformity described in the radioactive case are clear cases of states and events that are worse than life with a healthy bodily state.<sup>21</sup> Therefore, on the view I have proposed, they are harms.<sup>22</sup>

## 6. The happy child and reasons to benefit

In this section, I will discuss a puzzle about a claim called “the Asymmetry” and use this discussion to argue for claim (3) above, that there are reasons in favor of a course of action in virtue of the benefits to the future individuals who will not exist if the course of action is not taken.

The Asymmetry: There are reasons not to create someone who would have a life that would not be worth living; but there are no reasons to create someone who would have a life that would be worth living.

Many people find the Asymmetry plausible. Its first half is uncontroversial. Its second half appears to be true, in light of the following case:

Happy Child: A couple could have conceived, given birth to, and raised a child who would have had a happy life. They did not.

This couple does not do anything wrong. They are not blamable, and they do not need any justification for failing to conceive. They certainly should not feel badly about what they have done. Recognition of these facts leads people to deny that there are any moral reasons, in favor of conception, that the couple acts against.<sup>23</sup>

Nevertheless, I will now argue that the second half of the Asymmetry is false. I have endorsed:

- (4) An action harms a person if the action causes pain, early death, bodily damage, or deformity to her, even if she would not have existed if the action had not been performed.

From this claim, and the additional claim that whenever an action harms, there is thereby a reason against it, it follows that:

- (5) There are reasons against a course of action in virtue of the harms that will be experienced by the future individuals who will not exist if the course of action is not taken.

Besides following from (4), (5) is attractive in itself, in that it gives the right explanation of the truth of the first half of the Asymmetry: there are reasons against creation in virtue of the harms to the person who will be created, and these reasons are not outweighed.

Now consider an ordinary case of conception. In this case, it follows from (5) that there are some reasons against conception, due to the harms that any created person will inevitably suffer. Nevertheless, creation is permissible; what is the right explanation of its permissibility?<sup>24</sup> I claim that the right explanation is given by:

- (3) There are reasons in favor of a course of action in virtue of the benefits to the future individuals who will not exist if the course of action is not taken.

There are reasons in favor of creating a person stemming from the benefits she will have during her life; these reasons outweigh the reasons against creating. To put it simply, any couple who considers conceiving will worry about the harms that the child would inevitably suffer. In most cases, where the child is likely to have a good life, these worries need not keep them from conceiving—because they know that the child’s life will involve many wonderful features.

On my view, we do have some reasons to create every happy child we could create. But these are reasons *to benefit*; they stem from the way it would be non-relatively good for someone if we acted in a particular way. These reasons are very different from reasons *against harming*, which stem from the way it would be non-relatively bad for someone if we acted in a particular way. Reasons against harming have stronger force than reasons to benefit. We fail to benefit many people at every moment, in many different ways; and we need not feel badly about this. The couple who does not create a happy child does not do anything bad; they merely fail to do something that would be good.

## 7. Objection: the no regret argument

Consider this objection to my solution to the non-identity problem:

The No Regret Argument: The people affected by the policy do not (nor should they) regret that it was adopted. So they are not harmed by the policy and there is no reason against it in virtue of its effects on them.<sup>25</sup>

The objector asserts that, in order to have a legitimate complaint with an action, its victim must prefer the world *as it would have been if the action had not been performed* to the world *as it actually is*. The victim must regret that the action was performed, in that sense. But the people who are harmed by the policy would not have existed if the policy had not been adopted; they do not wish the policy had not been adopted.

In arguing that the No Regret Argument fails, I will also explain why the Worse-Off claim is fundamentally misguided. Two cases will illustrate that the following claim is true:

An individual can be harmed by an action and can have legitimate complaint about the action—indeed, she can be *impermissibly* harmed by the action—although she does not wish (nor should she wish) that the action had not occurred, and although the action makes her better off than she would otherwise be.

These cases are highly unusual. Most people who experience the harms described do not have stories like these. But what I need for my purposes is simply that these cases are *possible*, and they are.

**Rape:** A woman is raped, becomes pregnant, and ends up raising the child. The woman is remarkably able to separate the trauma of the rape from her attitude to the child, and they have a normal and healthy parent-child relationship. The woman's life is better, due to the value to her of the relationship with her child, than it would have been if she had not been raped, even taking into account the trauma of the rape. This woman loves her child. She does not wish that she had not been raped, because if she had not been raped, then her child would not exist.

**Nazi Prisoner:** A man was imprisoned in a Nazi concentration camp, where he suffered many harms. But his experience in the camp enriched his character and deepened his understanding of life, such that overall his life was better than it would have been had he not been imprisoned in the camp. He does not wish that the Nazis had not imprisoned him, because he so values what he gained from this experience.<sup>26,27</sup>

Both of these cases demonstrate that an action can impermissibly harm someone though it makes her better off than she would otherwise be, and though she does not wish the action had not been performed.

Someone might ask how these cases differ from the Surgery case, which is also a case in which a single action both harms and makes someone better off than she would otherwise be. There are two important differences. First, the

harms in the Rape and Nazi cases are such awful, gruesome harms that it is much harder for other factors to render them permissible. Second, the consideration available to outweigh the harm in Surgery is *the threat of worse harm of the same type*. The considerations available to outweigh the harms in Rape and Nazi are substantive benefits that come along with the harms. It is much easier for the threat of worse harm to render harm permissible, than it is for accompanying benefits to do so.<sup>28</sup>

These cases show that the Worse-Off claim is fundamentally misguided. It might have seemed that the Worse-Off claim only fails in non-identity cases, where an action affects whether the harmed person exists. These cases show that the Worse-Off claim fails even when restricted to cases where the harmed person exists independently of the action.<sup>29</sup>

These cases also demonstrate that there are two importantly different ways that harms and benefits can interact: there are two different senses in which benefits can *outweigh* harms. The benefits in these cases do outweigh the harms in that they are more beneficial than the harms are harmful: the total package of benefits plus harms leaves the person better off than he or she would otherwise be. But the benefits do not outweigh the harms in that they do not render it permissible to cause the harms. Another way of putting this is that the *reasons* to benefit do not outweigh the *reasons* against harm, though the benefits themselves outweigh the harms.<sup>30</sup>

I will now consider two objections to my discussion of these two cases. The first objection is that what makes the actions in each case impermissible is not *harm* at all, but rather the lack of *consent* of the parties to the actions. Raping someone and imprisoning someone, by their natures, involve lack of consent. The objector claims that the actions are wrong, despite the victims' lack of regret and despite the victims' being made better off, simply because there was no consent.

I agree with the objector that lack of consent plays a crucial explanatory role in both cases. But I affirm that *harm* also plays a crucial explanatory role. Part of the explanation of why it is wrong to rape, and why it is wrong to imprison someone as the Nazis did, is how awful it is for the victim to go through, and to have gone through, these experiences. Part of what makes these experiences awful to go through is that were not consented to. I claim that consent does play a crucial explanatory role, but that the harm to someone of going through these experiences also plays an important role in explaining what is wrong with these experiences. Furthermore, this harm provides reasons against the actions even in cases like the two I described above.<sup>31</sup>

The second objection is that my discussion of the cases trades unfairly on the fact that neither the Nazis nor the rapist either *foresaw* or *intended* the good effects of their actions. The objector claims that the agents' actions were wrong because, for all the agents knew, they were making the victims worse off.

This objection appears to be making the following proposal: the actions were *subjectively wrong*, impermissible given what the agent thinks the facts are,

but they were not *objectively wrong*, impermissible given the facts. This claim requires that if the agents had known what benefits would ensue, their actions would have been permissible. But suppose, improbably, that the prisoners or the rapist had a rare insight into the causal consequences that would occur, and acted out of a perverse paternalistic desire to make the victims' lives better. Their actions would still have been impermissible.<sup>32</sup>

## **8. Objection: can harm explain wrongness?**

One might object that I have not really explained why the actions in non-identity cases are wrong. Derek Parfit has argued that there is a general problem for anyone who tries to explain the wrongness of the actions in terms of harm, which is brought out in a comparison between a typical non-identity case and a variant of it. Consider these two cases:

Temporary Condition: Tammy has a temporary condition that will cause any baby she conceives now to be born deaf. She conceives now. If she had not conceived now, her condition would have cleared up in two months and she would have conceived a different baby who would not have been disabled.

Permanent Condition: Patty has a permanent condition such that if she ever conceives a child, the child will be deaf. She conceives.

To see Parfit's worry, we must make three assumptions about these cases, all of which I think are warranted. First, the state of being deaf is worse than a healthy bodily state, so causing deafness is harming someone.<sup>33</sup> Second, Tammy's action is impermissible. Third, Patty's action is less bad than Tammy's action. Tammy could have avoided causing anyone to be deaf by waiting two months to conceive; the inconvenience of waiting two months does not outweigh the serious harm she causes. Patty, on the other hand, must undergo the much more serious burden of not conceiving a child at all if she is to avoid causing someone to be deaf. Patty's reasons against conceiving are much stronger than Tammy's reasons against conceiving now, so Patty's action is less bad.

The worry is that my account cannot explain the difference between the two cases. Both Tammy and Patty harm their children, on my view. They harm their children in exactly the same way, by conceiving while they have a condition that causes deafness. The problem is that if Tammy's action is wrong in virtue of the harm to her child, Patty's action ought to be equally wrong in virtue of the harm to her child. Parfit glosses the problem as follows. He says that Tammy's action is not wrong merely in virtue of the harm to her child; rather, it is wrong because she had an alternative that would have involved no harm. He says this shows

that we cannot explain the wrongness of Tammy's action simply by appealing to the harm to her child.<sup>34</sup>

Parfit's comments bring up the question of what counts as *being the explanation* of wrongness. On my view, Tammy's action is wrong because she harms her child. But the explanation does not stop there, as Parfit points out. The harm to Tammy's child provides a reason against her action, *and the reasons in favor of her action are insufficient to render the harm permissible*: the benefits to the child cannot render the harm permissible because she has an alternative of providing parallel benefits without parallel harm, and the burden on Tammy herself of having to wait two months is insufficient to render the harm permissible. Does the fact that we must add this further information to our explanation show, as Parfit thinks it does, that the harm does not explain the wrongness of the action? I don't think it does. Harm *never* fully explains the wrongness of any action, in that sense. Perhaps some harms are such that nothing could justify them. But many harms are such that some circumstances could make them permissible. When we say that an action is wrong in virtue of the harm it causes, we are saying that there were no reasons in favor of the action that outweighed the reason provided by the harm.

Tammy's action harms her child, and there is a very weak reason for her to perform the action. Patty's action harms her child, but there is a much stronger reason for her to perform her action—it is the only way for her to conceive and gestate her own child. Thus, my view can explain the difference between the cases.

## 9. Actions that affect how many exist

In this section, I will expand on my solution to the Non-Identity Problem so that it can handle a variant of the Radioactive Waste Policy case. In discussing the case initially, I assumed that *how many* people would exist in the future is not affected by whether we adopt the policy, though *who* exists is affected: this enabled me to say that the benefits that would occur given either course of action are largely parallel. However, if we consider a variant of the case in which no matter how we act the population will grow, but it will grow *more* if we adopt the proposed policy, then the case might appear to be trickier. Adopting the policy will *harm* in ways failing to adopt it will not harm, but it will also *benefit* in ways failing to adopt it will not benefit, because more people will exist, and so more benefits will be enjoyed.

This case appears to raise a problem for my account. Suppose that the number of people who will exist if we do adopt the policy is  $x$ , the number who will exist if we do not adopt the policy is  $y$ , and  $x$  is substantially greater than  $y$ . Suppose further that the number of people who will be harmed if we adopt the policy is  $z$ , and that  $z$  is less than  $x$  minus  $y$ . (Plausibly this does not change our intuitions about the case—indeed the case originally needed only that *some* future people will be harmed. Even if only a small population within the future

population will suffer from the radioactive waste, that they will suffer and that we have an alternative in which we do not cause anyone to suffer in that way is enough to make adopting the Policy wrong.) Then the situation appears to be as follows:

Reasons in favor of adopting the Policy:

- if the Policy is adopted,  $x$  future individuals who will thereby exist will enjoy many benefits throughout their lives
- if the Policy is not adopted, we will continue to be somewhat inconvenienced

Reasons in favor of *not adopting* the Policy:

- if the Policy is not adopted,  $y$  individuals who will thereby exist will enjoy many benefits throughout their lives
- if the Policy is adopted,  $z$  individuals who will thereby exist will be harmed

The explanation I gave for the first version of the Radioactive Waste Policy case will not work here. Indeed, it might appear that the moral reasons interact as follows. The reasons to benefit are not parallel, so we are left with reasons in favor of adopting the Policy given by benefits to  $x-y$  persons. These reasons in favor of benefit, together with the reason against inconveniencing us with the existing policy, might well be stronger than the reasons against harming  $z$  individuals. After all,  $x-y$  is a greater number of persons than  $z$ ; and furthermore, the benefits to any one person who will exist in the future (whether that person is among the harmed or not) are substantially greater than the harms to any harmed person (on the assumption that the harms do not render lives not worth living). So, since adopting the policy provides certain benefits, not provided in the alternative, that are *more beneficial* than the harms it causes are *harmful*, it seems that the moral reasons to do with harm and benefit now tell in favor of adopting the policy.

That would be a devastating result. However, it does not follow from the view I am advocating. In section 7, I pointed out that harms and benefits can interact in two different ways. A group of benefits can outweigh a group of harms in that those benefits are more beneficial than the harms are harmful, yet nevertheless those benefits do not *morally* outweigh the harms: the reasons against benefit do not succeed in outweighing the reasons against harm. I now claim that this phenomenon is also at work in the revised version of the Radioactive Waste Policy case: the reasons against causing the harm are such *serious moral reasons* that the mere fact that some benefits come along with them, even that these benefits are not provided in the alternative course of action, is not enough to render the harms permissible.

This raises the question of when and why it is ever permissible to cause harm by creating (or by performing actions on which creations are counterfactually dependent). There are at least two ways that this can be permissible: either the alternative course of action is in some way tragic so there are compelling reasons to avoid it, or the harms in question are not so serious and there is no alternative of creating someone who is not similarly harmed.

An instance of the first sort of permission is given by a third version of the Radioactive Waste Policy case. Suppose we know that inevitably, several generations from now people will begin to suffer the kinds of harms that future individuals will suffer in the Radioactive Waste Policy case. We now have a choice, let's suppose, whether to continue to procreate or to simply let the human population die out without suffering those harms. The possibility in which humans cease to exist is morally serious enough, even considered simply as a loss to us and to the last generation of living humans, that it justifies creating and thereby causing the ensuing harms. In the second version of the Radioactive Waste Policy case, the alternative of pursuing our safe policy is in no way tragic: fewer people will exist in the future than would have had we adopted the policy, and we will continue to be inconvenienced, but neither of these facts is remotely tragic. Whether it is permissible for Patty to conceive in Permanent Condition, I claim, depends on how tragic her alternative is. Many people think it is impermissible for Patty to conceive; this can be explained by pointing out that the loss *to her* of getting to conceive a child is not a serious enough loss. In a variant of the case, suppose Patty is in a situation in which her only chance of getting to raise a child is by conceiving one, because adoption or embryo donation is unavailable. In this case, the loss to Patty is more serious. Whether we think her action is permissible should depend, I claim, on how tragic we think the deprivation of getting to raise a child is, together with how serious a harm it is to be deaf: if the harm is serious enough, we may think that no loss to Patty could be tragic enough to justify it.

The second kind of permission arises in ordinary creation. These are cases in which we harm those we create in creating them, but the harms in question are not so serious and there is no alternative of creating someone who is not similarly harmed. In an ordinary case of creation, we have no reason to think that the created child will suffer harms as serious as those in the Radioactive Waste Policy case, the Teenage Mother case, or the Temporary Condition case: most people live long relatively healthy lives. So the harms that our child will suffer are easier to render permissible than the harms in the cases that raise the non-identity problem. Also, we have no alternative of creating without causing similar harms: any person we create will inevitably be harmed in many ways.

It is sometimes worried that any solution to the non-identity problem will inevitably go too far, and require us not just to create those who will not suffer the harms in the cases that raise the problem, but also require us to create *the most perfect children we can create*. Because I have advocated a significant



difference between reasons against harm and reasons in favor of benefit, my view is not susceptible to this worry. While I claim that reasons against harm are serious reasons not easily outweighed, I do not claim the same thing about reasons to benefit. (I understand harm here as the causing of bad states and benefit here as the causing of good states. While the alleviation of harm is sometimes called “benefit” and the removal of benefit sometimes called “harm”, I am setting such usage aside.) My view does, however, have the consequence that in creating we have strong moral reasons to create, not *the most happy, most benefited, or most perfect* child we can create, but *the least seriously harmed* child we can create. This is a consequence I embrace. My view does not imply, however, that we should create the least seriously harmed child at any cost: we need not create a child who will live in a boring protected bubble over a child who will see the world and suffer more harms along the way. While reasons against harm are serious, they *can* be outweighed.

## **10. Objection: super-conservation**

A final objection goes as follows. The objector says that any correct explanation of why we ought not to adopt the permissive policy in Radioactive Waste Policy should also explain why we ought to adopt the Super-Conservation Policy in this case:

Super-Conservation: We have a choice whether to adopt a Super-Conservation Policy. If we do not adopt the Policy, we will continue with our reasonably good conservation policies, which will ensure that future generations enjoy a high level of well-being, and which allow our own level of well-being to rise steadily. If we do adopt the Policy, our own level of well-being will still rise, but slightly more slowly; future generations will enjoy a vastly higher level of well-being than they would otherwise have. Implementation of the Policy would have such widespread effects that different people will exist in the future depending on whether we adopt it. The benefits to future generations are much greater than the costs to us.

Parfit takes it to be clear (about a case like Super-Conservation, with different details) that we ought to adopt Super-Conservation. I think this is a mistake. We have no *obligation* to adopt Super-Conservation, though it would be very nice of us to do so.<sup>35</sup> I am not sure why Parfit takes it to be clear that we should adopt Super-Conservation. Perhaps he endorses:

- (6) In general, whenever one can undergo a small sacrifice to provide a significant benefit to many other people, one ought to do so, even if those other people are better off than oneself.

But surely (6) is false. Suppose that if poor children would each give up a delicious ice cream cone, then fifty happy and rich adults would receive massages every day for the rest of their lives. It would be nice of the children to give up the ice cream, but it would not be morally required. Our situation in Super-Conservation is similar, in that we are less well-off than those we could benefit.

### 11. Argument for the good method: is it wrong if we don't do it?

So far, in discussing the Radioactive Waste Policy case, I have provided an explanation of this claim:

If we do adopt the permissive radioactive waste policy, we do something wrong.

But this claim is also true:

*If we do not adopt the permissive radioactive waste policy, then: if we had adopted the permissive radioactive waste policy, we would have been doing something wrong; that is, there are reasons against adopting the policy.*

There are reasons against adopting the radioactive waste policy regardless of whether we do in fact adopt the policy. It is wrong regardless whether we adopt it. Why think there is any problem for saying this? There would be a problem if we adopted the following attractive, but bad, method:

Bad Method of finding out whether there are any reasons against an action due to harm:

First ask: Which things have moral status?

Then ask: If the action were performed, would it harm any of these things?

In answering the first question, we would list all actual persons, including past and future persons.<sup>36</sup> On my view (developed elsewhere)<sup>37</sup>, we would list many actual animals. Now suppose we are in the world in which we do not adopt the permissive policy. We then ask the second question: would adopting the policy have harmed any of those things? It would not have. Adopting the policy would have harmed some people, but those people do not actually exist; they are not on our list.

Some people worry about this problem. They worry that all actual reasons against harming must stem from actual things that actually have moral status. But this is a mistake. Reasons against a course of action need not rely on what the world is actually like; they can rely on what the world *would have been like* if the action had been performed. We should use this method instead:

Good Method of finding out whether there are any reasons against an action due to harm:

First ask : If the action were performed, would it harm any things?

Then ask: If the action were performed, would the harmed things have moral status?<sup>38</sup>

Again, suppose we are in the world in which we do not adopt the permissive policy. The first question asks: if we adopt the permissive policy, would it harm any things? The answer lists all those who would suffer as a result of the radioactive waste. The second question asks whether, if we had adopted the permissive policy, those things would have had moral status. They would have been persons, so they would have had moral status. This is the right result: there are reasons against adopting the permissive policy, due to harm, because if the policy had been adopted, it would have harmed things with moral status.

Consideration of the non-identity problem demonstrates that the Good Method rather than the Bad Method should be endorsed. This is related to work I have done elsewhere, in which I argue that early fetuses have their moral status contingently.<sup>39</sup> It then becomes important that we look at what the moral status facts *would have been* if a particular course of action had been performed; looking at the moral status facts as they actually are can give the wrong result.

## 12. Conclusion

We do harm and benefit those we create in creating them (or in performing actions on which creations are counterfactually dependent). I have developed a view that can explain the impermissibility of the actions in non-identity cases. I deny a purported *necessary* condition for harm:

Worse-Off: An action harms a person only if it makes the person worse off than she would otherwise have been if the action had not been performed.

I endorse a *sufficient* condition for harm:

- (4) An action harms a person if the action causes pain, early death, bodily damage, or deformity to her, even if she would not have existed if the action had not been performed.

I also endorse the following claim about *benefit*:

- (3) There are reasons in favor of a course of action in virtue of the benefits to the future individuals who will not exist if the course of action is not taken.

More generally, my view is that there is a reason against any action that would harm a person, and there is a reason in favor of any action that would benefit a person. These *individualistic reasons* can explain the moral facts in non-identity cases. Reasons against harm are *morally serious* reasons that are difficult to outweigh; the mere presence of benefits more beneficial than the harms are harmful is not sufficient to render harming permissible.

Finally, I advocate:

Good Method of finding out whether there are any reasons against an action due to harm:

First ask : If the action were performed, would it harm any things?

Then ask: If the action were performed, would the harmed things have moral status?<sup>40</sup>

## Notes

1. This problem is extensively discussed in Part IV of Parfit 1984. I believe it was introduced in Kavka 1982 and Parfit 1982.
2. Those who pose the problem assume that the following counterfactual is true: if a person's parents had not conceived together around the time they did, then that person would not have existed. This claim is merely a counterfactual about the nearby worlds in which the antecedent holds. It does not require a strong origin essentialist claim, because there may still be some remote possible worlds in which the person's parents do not conceive together around the time they did, and yet that person still comes to exist.
3. Derek Parfit draws a lesson like this. Gregory Kavka argues "that individualistic principles of obligation...cannot be relied on to yield correct results when applied to cases involving the creation of persons," Kavka 1982, 103.
4. The question at issue is whether a particular action results in harm in a way that is morally relevant (that is, the action results in harm and there is thereby a reason against the action). Sometimes someone *causes* a can to be squashed without *squashing* the can herself. Similarly, some people think there is an important distinction between *causing harm*—even where this is relevant to the morality of the action—and *harming*; I am setting this distinction aside. Whenever an action causes harm and there is thereby a reason against the action, I say that the action harms.
5. The second goal also has further significance. This paper is part of a larger project, in which I argue for an answer to the question "Which things have moral

status?” I understand one of the central aspects of something’s having moral status to be as follows: the mere fact that an actions *harms* something with moral status provides a reason against the action. The interest of the question “Which things have moral status?” thus partly depends on the extent to which reasons against harm are importantly morally explanatory. Achieving the second goal lays to rest the worry that the non-identity problem shows reasons against harm to be less morally important than they generally appear.

Parsons (2003) sees the non-identity problem as posing an objection to “the person-affecting view”—roughly, that right and wrong reduce to benefit and harm to particular persons.” But we may well be interested in vindicating a harm-based explanation of non-identity cases, which seem so clearly to warrant an explanation in terms of harm, without being inclined to the strong claim made by “the person-affecting view.”

6. The fourth goal also bolsters a view I’ve developed elsewhere about how to evaluate merely counterfactual harm. See the end of section 11.
7. Parfit 1984 mentions the surgery case; Parfit 1987 stresses the importance of the surgery case.
8. Reasons interact in many different ways. While it is sometimes useful to talk of one reason being stronger than another, it by no means follows that all reasons have strengths that can be compared, or that reasons simply add up with the stronger one winning out. But sometimes reasons are of the right type to be weighed against each other, and the stronger do win out. Here we have a reason against causing a harm to me, and another reasons against allowing serious harm to befall me: both are reasons to do with effects on a single person, and both involve possible *harm* to me—this enables the reasons to be weighed directly against each other, with the fact that one is much stronger than the other settling what may be done.
9. Woodward (1987, 807) rightly stresses the importance of consent to the permissibility of performing surgery, where consent is possible. But the surgery case raises a problem even if we assume that consent is granted. The existence of consent is not a sufficient justification for harming someone; it is merely a necessary part of many justifications.
10. Shiffrin (1999, 125–126) discusses a rescue case in which the rescuer must break the victim’s arm to save her life. She argues that the rescuer both harms the victim and prevents more serious harm to the victim, and that this is why the action is permissible.
11. This claim, that the benefits are more beneficial than the harms are harmful, is plausible. There is nothing in the description of the case to suggest that those harmed do not have lives worth living. Most people do, even those who suffer substantial harms.
12. In endorsing (4) and denying the Worse-Off claim, I do not deny that we sometimes use the word “harm” to mean “all-things-considered harm.” Indeed, the phrase “First, do no harm”—commonly attributed to the doctors’ Hippocratic Oath—uses the word in this way. The important point is that there can be harm even when there is not all-things-considered harm, and there can be a reason against an action in virtue of harm even when there is not all-things-considered harm.

13. A variant of this case is described in Parfit 1984.
14. A variant of this case is described in Parfit 1975.
15. Schwartz (1978) denies that the teenage mother has any reason not to harm the fetus. Rather, he says, the teenager has a reason not to cause herself the burdens of very young motherhood; and that is the only reason against her action (absent further details). While Schwatz is right that the burdens on herself provide reasons against the teenager's action, his account seems to miss the primary reason against conceiving.
16. Bayles (1976) argues that there cannot be an account of these cases based on individualistic considerations. He argues for the following impersonal principle: "There is good reason for legislation to prevent the birth of persons who would lack substantial capacity to achieve or take advantage of a quality of life of level  $n$  or whose existence would decrease the number of people who might live with a quality of life at that level" (302). The principle could be modified from a legal claim to a moral claim. One worry with such a solution is that the teenager's action might well be wrong even if her child would not end up so seriously poorly off as to be below whatever cut-off line would be correct.
17. There is an important respect in which the Teenage Mother and Temporary Condition cases are different from the Radioactive Waste Policy case. In the two conception cases, I stipulated that if the woman had not conceived now, she would have conceived later. But her action would be equally wrong if it were unclear whether she would conceive a healthy child later; what matters is that she has the opportunity. My account may appear to be unable to accommodate this fact, because I rely on the claim that there are parallel reasons, due to the benefits to the created child, in favor of conceiving now and *not conceiving now*. I think that we can still say that there are reasons in favor of not conceiving now, even when it is not certain or not true that the mother would ever conceive later, in virtue of the benefits she *could* bestow later if she does not conceive now.
18. I do not claim that the point of comparison only involves a healthy bodily state; I claim that it also involves a healthy mental state. This explains why the harms to the child in Teenage Mother do count as harms.
19. Woodward (1986) explains that his claim that harm occurs to the future individuals (in a case like Radioactive Waste Policy) appeals to a comparison with an "unattainable" situation in which these people exist and are not harmed in the way they are actually harmed (817). Similarly, Hanser (1990, footnote 6), says that we can call an action bad for someone if it would have harmed them—by which he appears to mean, if it would have made the person worse off—if the situation had been similar except that we were not affecting who exists. These accounts offer an alternative idea of what the point of comparison might be. The accounts might be open to the complaint that they beg the question against those who think that Non-Identity Problem cases are importantly different from similar cases in which our actions do not affect who exists.
20. Hanser (1990, 65) puts a similar point as follows. He considers a case in which our negligent policies cause someone to be born with a genetic predisposition to die young. He says that having a genetic predisposition to die young is bad for someone. If we are responsible for the cause of the genetic predisposition, then we are responsible for the badness to the person. His claim that having such a

- genetic disposition is bad for someone requires either a notion of non-relative badness or a baseline of a healthy life for comparison.
21. It is important that causing someone to be in a state which is *in any respect* worse than a healthy bodily state is sufficient for harming.
  22. Shiffrin (1999, 123) advocates a view of harm and benefit according to which we identify harms by “a list of evils like broken limbs, disabilities, episodes of pain, significant losses, death” and benefits by another list, “material enhancement, sensual pleasure, goal-fulfillment, nonessential knowledge, competitive advantage.” She also urges that such a view provide some explanation of what unifies members of each list. I have proposed that my list of harms is unified by comparison with a healthy state (though I haven’t claimed that all harms meet this condition). Shiffrin proposes, roughly, that all harms interfere with a person’s ability to control her life through her will.
  23. Utilitarians, who claim that we ought to do whatever would result in a greater sum of happiness, would deny the Asymmetry.
  24. Thanks to Carolina Sartorio for pressing this question.
  25. Parfit (1984, 373) offers this argument.
  26. The prisoner need not be callous: he may also wish that the Nazis had not imprisoned anyone. He simply prefers the world as it is to the world as it would have been if he had not been imprisoned; he could prefer the world as it would have been if the Nazis had not imprisoned anyone, to both these worlds.
  27. Woodward (1986, 809) describes a case like this in criticizing a version of the Worse-Off claim; he does not mention its relevance to the No Regret argument.
  28. Shiffrin (1999) points out that this is an important distinction, and that the mere fact that greater benefits come along with harms does not render causing the harms permissible.
  29. Other cases would show that lack of regret can come apart from whether someone is worse off. In the Rape and Nazi cases, the victims do not regret that the action occurred, and the action also makes them better off than they would otherwise have been. But a victim might not regret that an action occurred, even though she is made worse off by the action. Someone forced to carry a teen pregnancy to term might well not wish she did not have the child she in fact raised and loves, while knowing that she would have been better off if she had been allowed to abort—she nevertheless prefers the world as it is to the world in which she would have been better off but would not have had the particular child she has come to love.
  30. Woodward 1986 makes two points about the No Regret argument that are worth discussing. First, Woodward says that those harmed could *regret* an action in virtue of its having been *morally wrong*, even if the action resulted in their being better off than they would otherwise be (823–824). I think that this response simply uses the word “regret” in a different way than is intended by the proponent of the No Regret argument: the issue is whether those harmed do or should wish the harming action had not been performed. Second, Woodward says that, given the fact that I do not wish a state of affairs *q* did not obtain, it does not follow that I do not wish a state of affairs *p*, that was necessary for *q*, did not obtain. He gives the example that I may not regret making amends for a wrong I did (*q*), while I do regret committing the wrong (*p*), even though without

it the amends would have been impossible. I think this example is inadequate to do the work we need in these cases. I can have the attitudes Woodward describes because I prefer the world without either wrong or amends to the world as it is, which I in turn prefer to the world with the wrong but without the amends. The situation in the two cases I described is different. The victims of the harm actually prefer the world with the harm and resulting benefits to the world without the harm. I think Woodward's move does suggest an interesting point: we can wish for things that are impossible. The raped woman could wish that she had gotten this very child without conceiving with the rapist. The Nazi prisoner could wish that he had gained the very same insights without suffering. But I am not convinced that the possibility of such wishes affects the claims made in the No Regret argument.

31. It's compatible with what I have said here, that there can sometimes be impermissible rape involving no bad experiences, and perhaps also impermissible concentration camp imprisonment involving no bad experiences, where the impermissibility is solely located in the violation of consent. While this is possible, it is not usual, and such cases have importantly different moral features than the two cases I described.
32. Woodward (1986) points out that we do not generally think that intentions matter to wrongness: "the usual view is that whether or not the benefit is intended or 'aimed at' (Parfit's phrase) makes very little difference to the justifiability of the action, although it may perhaps make a difference to the blameworthiness of the action if we decide that his action was unjustifiable" (footnote 8). He says, about a case similar to the Nazi Prisoner case, that even if the Nazis intended the benefits to the prisoner, and acted for his own sake, their actions would be impermissible. He also claims that the fact that a policy "happens to have, on balance, consequences such that it leaves [a harmed person] no worse off than he would be under any alternative policy is simply irrelevant" (814). This suggests that benefits can never make any harm permissible; I think this is too strong. It is an interesting question why the Nazis' behavior to the prisoner, no matter their motive, is impermissible, whereas it is sometimes permissible to force soldiers to go through brutal boot camps, or to make children endure apparently harsh parenting. Those are cases in which the benefits to the individual appear to render the harms permissible.
33. The problem can also be developed for cases of conception that lead to suffering, rather than deafness.
34. In Parfit 1984; the argument is repeated in Parfit 1986 (861).
35. Some writers who discuss the Non-Identity Problem agree with me on this point. See Hanser 1990 (66). Woodward 1986 (820) goes even further, saying there may be a moral objection to adopting a policy like Super-Conservation.
36. In saying that we would list past, present, and future persons, I am saying that the first question should be read tenselessly. Some Presentists will say that even a tenseless reading of that question will yield only presently-existing persons. This raises a temporal puzzle analogous to the modal puzzle I explain in this section; it is amenable to a parallel solution to the one I provide.



37. On my view, something has moral status throughout its life just in case there is a time in its life at which it is conscious. This view is developed in my 1999 and 2003.
38. Woodward 1986 offers essentially the same response to this problem. He understands the problem as the complaint that we cannot “appeal to the rights” of people who do not and will not exist. He says that “a certain course of action would be wrong...[if] it would involve the creation of rights and obligations that would probably inevitably be violated” (821).
39. I argue in my 1999 that whether a pre-conscious fetus has moral status depends on whether it is actually going to develop into a person.
40. I would like to thank the following people for helpful discussion and comments: Dan Brock, Alex Byrne, Joshua Cohen, Austin Dacey, Alex Guerrero, Peter Graham, Frances Howard-Snyder, Sarah McGrath, Michael Rescorla, Carolina Sartorio, Thomas Scanlon, Brad Skow, Paul Tudico, and Ralph Wedgwood. I am also grateful to audiences at Brandeis University, the Bellingham Summer Philosophy Conference, Brown University, the Eastern APA Meeting, the Harvard Moral and Political Workshop, the LSU Symposium on Theoretical and Applied Ethics, the MIT M.A.T.T.I. Group, New York University, the University of British Columbia at Vancouver, the University of Calgary, and the University of Toronto.

## References

- Bayles, Michael D. 1976 “Harm to the Unconceived” *Philosophy and Public Affairs* 5: 3: 292–304.
- Hanser, Matthew 1990 “Harming Future People” *Philosophy and Public Affairs* 19: 1: 47–70.
- Harman, Elizabeth 1999 “Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion” *Philosophy and Public Affairs* Fall 1999, 28: 310–324.
- Harman, Elizabeth 2003 “The Potentiality Problem” *Philosophical Studies* May 2003, 114: 1–2: 173–198.
- Kavka, Gregory S. 1982 “The Paradox of Future Individuals” *Philosophy and Public Affairs* 11: 2: 93–112.
- Parfit, Derek 1986 “Comments” *Ethics* 86: 832–872.
- Parfit, Derek 1982 “Future Generations: Further Problems” *Philosophy and Public Affairs* 11: 2: 113–172.
- Parfit, Derek 1984 *Reasons and Persons* Oxford: Clarendon Press.
- Parfit, Derek 1987 “A Reply to Sterba” *Philosophy and Public Affairs* 16: 2: 193–194.
- Parsons, Josh 2003 “Why the handicapped child case is hard” *Philosophical Studies* 112: 2.
- Schwartz, Thomas 1978 “Obligations to Posterity” in Richard Sikora and Brian Barry, eds., *Obligations to Future Generations* Philadelphia: Temple University Press, 3–13.
- Shiffrin, Seana 1999 “Wrongful life, Procreative Responsibility, and the Significance of Harm” *Legal Theory* 5: 2: 117–148.
- Woodward, James 1986 “The Non-identity Problem” *Ethics* 96: 4: 804–831.
- Woodward, James 1987 “Reply to Parfit” *Ethics* 97: 4: 800–816.