

HOW TO SOLVE THE NON-IDENTITY PROBLEM

David Boonin

The purpose of this paper is three-fold: to explain what the non-identity problem is, to critically examine the various solutions to the problem that have been offered, and to propose an alternative solution.

I. THE PROBLEM

The non-identity problem arises from the fact that in some situations, our choices determine not only the quality of life that people will enjoy in the future, but the identity of the people who will exist to enjoy it. In such cases, the application of moral principles that seem relatively uncontroversial in more familiar contexts can yield results that most people find unacceptable. The problem is thus one of determining how, if at all, these common sense moral beliefs must be revised in light of their implications for our obligations to those future people whose identity we can (at least in part) control. The problem seems to have been discovered independently in the late 1970's by Derek Parfit, Thomas Schwartz, and Robert M. Adams, and is now most closely identified with Parfit, whose 1976 article "On Doing the Best for Our Children" was among the first to report it, whose seminal 1984 book *Reasons and Persons*, contains its fullest and most influential treatment, and who gave it the name by which it is now most commonly known.¹

1.1. Two Cases

Since the problem is one of relating general principles to specific instances, it is best approached by means of its application to a particular set of cases. I will therefore begin with a pair of cases that help to bring out the nature of the problem.

Betty: Betty takes her newborn baby for a checkup. The doctor says that there is some good news and some bad news. The bad news is that, as things now stand, the baby is going to develop a handicap. The doctor explains that the handicap will be significant, meaning something that uncontroversially

diminishes one's quality of life in a non-trivial way (e.g., more like blindness than like color-blindness). It will be non-terrible, meaning that although life with this handicap is considerably worse than life without it, it is nonetheless clear that it does not come close to making life worse than no life at all (assuming that such a thing is possible). And it will be irreversible, meaning that once the handicap develops, there will be nothing that anyone can do to treat it. The good news, though, is that Betty can prevent all of this from occurring simply by giving the baby a tiny pill once a day for the next two months. The pill is easy to administer, has no side-effects, and will be paid for by Betty's insurance company. Fully understanding all of this, Betty decides that having to give the baby a pill once a day for two months is too inconvenient and so chooses to throw the pills away. As a result, she ends up with an incurably blind child rather than a sighted child.

Wilma: Wilma is not yet pregnant, but is planning to try to have a baby and so goes to the doctor for a pre-conception checkup. The doctor says that there is some good news and some bad news. The bad news is that if Wilma conceives, as things now stand, she will conceive a child with a significant, non-terrible, irreversible handicap. The good news, though, is that Wilma can prevent this from occurring simply by taking a tiny pill once a day for the next two months before conceiving. The pill is easy to take, has no side-effects, and will be paid for by her insurance company. Fully understanding all of this, Wilma decides that having to take a pill once a day for two months before conceiving is too inconvenient and so chooses to throw the pills away and conceive at once. As a result, she ends up with an incurably blind child rather than a sighted child.

I assume that you agree that what Betty does is immoral. And I assume that you agree that what Wilma does is immoral as well. If you reject either or both of these assumptions, then the non-identity problem will not arise for you, and you may wish to stop reading at the end of this sentence. But most people do think that what Betty and Wilma do in these cases is quite wrong. Indeed, most people agree that what they do are equally wrong. The cases, after all, are identical but for the fact that Betty makes her choice after her child is born and Wilma makes her choice before her child is conceived,² and it is difficult to imagine how the mere timing of their choices could make a moral difference.

But there is a problem. The problem is that there is a seemingly sound argument that apparently demonstrates that there is a morally relevant difference between what Betty does and what Wilma does. Indeed, the argument seems to establish that while what Betty does is quite wrong, what Wilma does is not even a little bit wrong. The argument rests on premises that most people seem to accept, yet produces a conclusion that virtually everyone rejects. This is what gives rise to the non-identity problem.

1.2 The Argument

The argument begins by pointing to a difference between the two cases that Parfit refers to as the difference between “same person” and “same number” cases. In the case of Betty, she will have one and the same child regardless of her choice. Suppose that the child she has just given birth to is named Bam-Bam. In that case, regardless of whether Betty gives her baby the pill once a day for two months, Bam-Bam will exist. If she gives him the pill he will grow up to be sighted and if she does not, he will grow up to be incurably blind, but either way, it is he who will grow up and be her son. The case of Betty is thus a “same person” case. But the case of Wilma is different. If Wilma takes a pill once a day for two months before she conceives, she will conceive a child who is different from the child she will conceive if she throws the pills away and conceives at once. This is because the sperm and egg that would come together two months from now would be different from the sperm and egg that would come together now and because the child’s identity is a function of the sperm and egg whose coming together result in its existence. For purposes of illustration, it may help to imagine that if Wilma conceives now she will have a girl and name her Pebbles, and that if she takes a pill once a day for two months before conceiving she will have a boy and name him Rocks. On this understanding, it is not the case that whatever choice Wilma makes, the same child will exist. Rather, either Pebbles will exist and be incurably blind, or Rocks will exist and will not be blind. Thus, regardless of Wilma’s choice she will have the same number of children, namely one, but they will not be the same person. So while the cases of Betty and Wilma both involve the same number choices, the case of Betty involves a same-person choice while the case of Wilma does not.

It is not immediately obvious that this difference makes a moral difference. But the argument that gives rise to the non-identity problem maintains that it does for the following reason: When Betty chooses to throw the pills away rather than give them to her son once a day for two months, her choice makes Bam-Bam much worse off than he would have been had she not done so. But when Wilma chooses to throw the pills away rather than take them once a day for two months before conceiving, her choice does not make Pebbles worse off than she would otherwise have been. After all, if Wilma instead waited two months before conceiving, then Pebbles would not have been conceived in the first place. Rocks would have been conceived instead. And since we are stipulating that the significant and irreversible handicap that Pebbles is born with is not terrible, it follows that the act of conceiving Pebbles does not cause her to live a life that is worse than no life at all (assuming that such a thing is possible), and so does not make her worse off than she would have been had Wilma not conceived her. I will call this

P1: Wilma’s act of conceiving now rather than taking a pill once a day for two months before conceiving does not make Pebbles worse off than she would otherwise have been.

This claim is the foundation of the argument that gives rise to the non-identity problem.

The argument conjoins to this claim what seems to be a very plausible analysis of what it is to harm someone, namely: if your act harms someone, then it makes that person worse off than they would have been had you not done the act. This analysis seems amply confirmed by the fact that if a person is accused of causing harm to someone, it is standardly taken as a sufficient rebuttal to the claim if they can establish that they have not made their alleged victim worse off than they would otherwise have been. So we have

P2: If P's act harms Q, then P's act makes Q worse off than Q would have been had P not done the act

From P1 and P2 we get:

C1: Wilma's act of conceiving now rather than taking a pill once a day for two months before conceiving does not harm Pebbles

The argument then adds a further stipulation: that Wilma's act does not harm anyone else, including Wilma herself. Now, of course, in many real life cases, the addition of a handicapped child to the world will probably impose a cost on someone or other. If nothing else, another blind child in the classroom is likely to require additional resources and thus to raise the total costs that must be borne by the taxpayers. But adding this stipulation to the argument seems perfectly reasonable nonetheless. After all, on the assumption that you do agree that what Wilma does is morally wrong, it is not as if you react to the case by thinking "oh, the poor taxpayers; what a terrible thing that Wilma has done to them". In deciding whether you think that Wilma has done something wrong, you do not first demand information about whether her choice will impose costs on third parties. Since it seems clear that your belief that Wilma has done something wrong is independent of whether her act harms anyone else, it seems clear that you believe that she has done wrong even if she has harmed no one else. It therefore seems fair to simplify the argument by simply stipulating that her act has not harmed anyone else and to see whether this will prevent us from concluding that her act was wrong. In short, we have reason to accept, even if only for the sake of the argument

P3: Wilma's act of conceiving now rather than taking a pill once a day for two months before conceiving does not harm anyone else

And from P3 and C1, it follows that

C2: Wilma's act of conceiving Pebbles does not harm anyone

At this stage of the argument, a basic moral principle is invoked, one that most people seem to accept. In its simplest form, it is the idea of "no harm, no foul," the thought that if an act harms no one, then the act is not wrong. For purposes of analysis, however, it may prove useful to break this claim down into two parts: the claim that if an act harms no one, then it wrongs no one, and the claim that if an act wrongs no one, then it is not morally wrong. The first part maintains that

if an act does not harm a particular person then that person has no moral claim against the act's being done. The second maintains that if an act is such that no particular person has a moral claim against its being done, then it is not wrong to do the act.

With this basic moral principle in mind, the argument concludes as follows. First, we set out the first part of the "no harm, no foul" principle:

P4: If P's act does not harm Q, then P's act does not wrong Q

From P4 along with C2 it follows that

C3: Wilma's act of conceiving Pebbles does not wrong anyone

The argument then adds the second part of the "no harm, no foul" principle:

P5: If P's act does not wrong anyone, then P's act is not wrong

From P5 and C3 we are then entitled to conclude that:

C4: Wilma's act of conceiving Pebbles is not morally wrong.

The premises seem right. The conclusion seems wrong. That's the problem.

1.3. The Conclusion

Before moving on to consider ways in which we might overcome the non-identity problem, I want to say a bit more about just what the problem is supposed to be. In particular, there are two different claim about Wilma's choice that an argument based on the fact of non-identity might attempt to force us to accept:

Claim 1: Wilma does not make the morally worse choice

Claim 2: Wilma does not make a morally wrong choice

In some contexts, these two claims might be coextensive. Suppose Wilma's choice is between killing a bunch of innocent people and not killing a bunch of innocent people. Then if she kills them, she makes the morally worse choice and does something morally wrong. But it is at least not obvious that every time someone makes the morally worse choice they make a morally wrong choice. If there are, as commonsense morality supposes, acts that are supererogatory, then there are acts that it is morally better to do than not to do but which are nonetheless not wrong not to do. A person who chooses not to do a supererogatory act rather than to do one, on this account, makes the morally worse choice, but does not make a morally wrong choice.

Since making the morally worse choice and making a morally wrong choice are not the same thing, it is important to be clear about what, precisely, the problematic claim is supposed to be that the non-identity problem is supposed to force us to accept. Parfit himself is not sufficiently clear on this point. In the context of one of the cases he discusses, for example, he says that the problem is that we think the potential mother in question "ought" to wait and conceive later (1984: 358). In other places, he says that we would have an "objection" to her choice not to wait

to conceive (e.g., 1984: 359). At still other points, he says that she would have a “moral reason” to wait. All of this is consistent with both possible formulations of the non-identity problem. If we think that Wilma is making a morally wrong choice, we will object, think that she has moral reason to do otherwise and think that she ought to do so. But if we think that Wilma is merely making the morally worse choice, we may feel the same way. Even if we think that her choice is not morally wrong, we might still think that she has moral reason not to make it and might object to her choice for that reason.

So perhaps Parfit does not think there is an important difference here. But I think this is a mistake. The non-identity problem poses a problem to the extent that it generates a conclusion that seems to be strongly at odds with common-sense beliefs. To the extent that common-sense beliefs oppose both claim 1 and claim 2, the claim that Wilma does not make the morally worse choice and the claim that Wilma does not make a morally wrong choice, an argument for either claim would suffice to pose a serious problem. But while the argument presented in section 1.2 makes clear how one can seem to be forced into accepting claim 2 by arguing only from premises that seem quite difficult to deny, there is no parallel argument for accepting claim 1. In order for an argument to convince us that claim 1 was true, we would have to replace the principle behind P4 and P5 with something quite different. Instead of saying that if the choice between one of two acts harms no one, then the act is not wrong, the argument would have to depend on the claim that if the choice between one of two acts harms no one, then it is not the morally worse choice. But this latter claim is not really a part of common-sense morality in the first place. If anything, common-sense morality clearly rejects it. Common-sense morality recognizes the category of supererogatory acts, and if such acts exist, then there are many cases in which you can choose between two acts, neither of which harm anyone, but one of which is still morally better than the other. And if one is morally better than the other, then the other must be morally worse. So it seems to me that the non-identity problem is only a problem when construed as an argument against the apparent *wrongness* of Wilma’s choice, not against its apparent *worseness*. For the purposes of the remainder of this paper, that is how it will be understood.

2. REJECTING A PREMISE

The argument that results in the non-identity problem contains five premises. It seems quite clear that if all five of the premises are true, then the conclusion must be true as well. It therefore seems that the only strategy available for solving the problem is to reject one of the premises. This is how virtually everyone who has tried to solve the non-identity problem has proceeded. I will argue in this section that none of these approaches succeeds.

2.1. *Rejecting P1*

P1 maintains that Wilma's act of conceiving now rather than taking the tiny pill once a day for two months before conceiving does not make Pebbles worse off than she would otherwise have been. This seems clearly to be true. If Wilma had not conceived when she did, after all, then Pebbles would not have existed at all. But there are two kinds of worry that can nonetheless be raised about P1.

The first worry arises from the suspicion that accepting P1 requires us to make a kind of comparison that many people believe to be impossible to make. For it can seem that in order to accept the claim made by P1, one must be able to compare how things are for Pebbles when she exists and is blind with how things are for Pebbles when she does not exist. And since there is nothing that it is like for Pebbles not to exist, it can seem that this comparison is unintelligible. If this comparison cannot be made intelligible, and if its being made intelligible is necessary in order for P1 to be true, then P1 cannot be true.

It is not clear to me if the comparison between how things are for Pebbles when she exists and is blind and how things are for Pebbles when she does not exist is intelligible or not. Nor is it at all clear to me that one need make such a comparison in order to endorse the truth of P1. P1, after all, does not purport to identify a particular relation that obtains between blind, existing Pebbles and non-existing Pebbles. Rather, it simply asserts that there is a particular relation that does *not* obtain between them, the relation of being worse for Pebbles. And whether this relation fails to obtain because it cannot obtain or simply because it happens not to obtain does not make an obvious difference to the truth value of P1 (it is not clear to me, for example, why we shouldn't say that the statements "three is not bigger than seven" and "green is not bigger than justice" are both true even though one number can be bigger than another while a color cannot be bigger than a virtue. It is clear that "justice is smaller than green" is false, but it is not clear that "green is not bigger than justice" is false).

But even if the comparison between existing and non-existing Pebbles really is unintelligible and even if the truth of P1 really does depend on its being intelligible, this will do nothing to solve the non-identity problem. For if the purported unintelligibility of the comparison between existence and non-existence prevents us from affirming P1's claim that Wilma's act does not make Pebbles worse off, then it will equally prevent us from affirming the denial of P1, the claim that Wilma's act does make Pebbles worse off. If we cannot claim that Wilma's act does make Pebbles worse off, then we cannot claim that Wilma's act harms Pebbles. And if we cannot claim that Wilma's act harms Pebbles, then we will still be stuck on the path that leads to the conclusion that generates the non-identity problem: if we cannot affirm that Wilma's act does harm Pebbles, then how can we explain why her act is wrong? Regardless of what we think about the intelligibility and relevance of the comparison between existence and non-existence, then, it remains the case that the claim "Wilma's act makes Pebbles

worse off than she would otherwise be” is false. And this is all that is required for the argument to proceed.

There is a second and very different worry that can be raised about P1. This second complaint maintains that I have unfairly misrepresented its content. The point of P1, after all, is to claim that Wilma’s act of conceiving now rather than taking the pills for two months before conceiving does not make her child worse off. To help to make the point more vivid, I have given the child a name so that we have a clear and specific reference point to appeal to. But, the critic charges, this substitution of terms is illicit. If we represent P1 as saying that Wilma’s act does not make “her child” worse off, rather than saying that it does not make “Pebbles” worse off, we will better represent the argument’s starting point. And if we do this, the critic maintains, we will also see that the claim is false.

The problem for P1, on this account, arises from the familiar *de re/de dicto* distinction. The claim that Wilma wants to marry the tallest man in New Haven, for example, could mean two different things: that there is some specific person that Wilma wants to marry and who happens to be the tallest person in New Haven (the *de re* reading) or that she wants to marry whatever person turns out to satisfy the description “the tallest man in New Haven” (the *de dicto* reading). Similarly, the claim that Wilma’s act does not make her child worse off is subject to two interpretations: the *de re* claim that her act does not make the specific child that she does conceive (Pebbles) worse off, and the *de dicto* claim that her act does not make whatever person turns out to satisfy the description “her child” worse off. But, the critic then argues, the *de dicto* interpretation of the claim is false, and, more importantly, its contrary is true. In the *de dicto* sense, Wilma’s act really does make her child worse off than it would otherwise have been. Her act makes it the case that “the child that she conceives” is blind rather than that “the child that she conceives” is sighted, and this is to make it the case that “the child that she conceives” is worse off in the *de dicto* sense.

This critic is surely correct to maintain that P1 is false on the *de dicto* reading. But, as with the first worry about P1, this concession can do nothing to solve the non-identity problem. The significance of P1, after all, is that it sets the stage for the further claim that since Wilma’s act does not harm her child, it does not wrong her child. And simply noting that P1 is true in the *de re* sense but false in the *de dicto* sense will do nothing to prevent us from reaching this further claim unless we are willing to assert that Wilma’s harming her child in only the *de dicto* sense could count as wronging her child (in any sense). But the claim that harming someone in only the *de dicto* sense could count as wronging someone is plainly implausible. If Wilma chose to adopt a blind child rather than a sighted child, for example, this choice would also harm “her child” in the *de dicto* sense, but it would be absurd to think that her act thereby wronged her child or anyone else. For purposes of the remainder of this paper, therefore, I will simply stipulate that when I say that Wilma’s act does not harm someone or does not make someone

worse off, this is meant in the *de re* sense, and that the same is true of the claims that if an act does not harm someone it does not wrong them and that if an act does not harm anyone then it is not wrong. As long as we keep this in mind, this further worry about P1 can also be set aside.

2.2. Rejecting P2

P2 maintains that if an act harms someone, then it makes them worse off than they would otherwise have been. If an act does not make someone worse off than they would otherwise have been, the act does not harm them. In order to reject P2, therefore, we would have to identify a morally relevant sense of harm in which an act could harm someone without making them worse off than they would otherwise have been. In “Harming Future People,” Matthew Hanser proposes to do just that.³ Hanser points out that there is a difference between the following two claims:⁴

P2: If P’s act harms Q, then P’s act makes Q worse off than Q would have been had P not done the act

P2': If P’s act harms Q, then P’s act makes it the case that Q is harmed more than he would have been harmed had P not done the act.

He demonstrates that these two principles are not coextensive by applying both of them to a case of Parfit’s called the Risky Policy.⁵ In the Risky Policy case, a society embraces an energy policy that predictably results, many centuries later, in the release of radiation that painlessly kills thousands of people. Because of the pervasive effects of choosing one energy policy over another, however, the people who are killed by the radiation would not have existed in the first place had the Risky Policy not been chosen. As Hanser points out, the people who are killed by the radiation are not harmed by the adoption of the Risky Policy according to P2. The adoption of the policy does not make them worse off than they would otherwise have been, since they still have worthwhile lives and would not have existed at all had the policy not been selected. But the adoption of the Risky Policy does harm those who are killed by the radiation according to P2'. Adoption of the policy causes the radiation to leak, which causes their deaths, which harms them. So there are harms that they suffer given the adoption of the policy. And since if the policy were not adopted, they would not exist at all, it follows that if the policy were not adopted, they would suffer no harms at all. The act of adopting the policy, therefore, does harm them in the sense that they are harmed more if the policy is adopted than if it is not adopted.

Hanser is clearly correct in pointing out that P2' represents an alternative analysis of harm in the morally relevant sense. And he is also clearly correct to say that this alternative account, if accepted, would eliminate the non-identity problem, at least in cases like that of the Risky Policy.⁶ The question, then, is whether we have any reason to prefer P2' to P2. Hanser’s argument for accepting

P2' rather than P2 runs as follows:

Parfit would admit that a Same People Choice analogue of the Risky Policy [i.e., one in which the same people would exist in the future regardless of which policy choice is made] that resulted in a similar catastrophe would harm people in the morally relevant sense. And it seems that the proper account of this would be as follows. First, the people who would die in the catastrophe would suffer harms. Many would be in the prime of life; death, then, would surely come to them as the greatest of evils. Second, we would be *responsible* for their deaths, since we would be responsible for the presence of the radioactive material. . . . But both parts of this explanation hold for our choice of Parfit's Risky Policy as well: nothing in the explanation turns on whether or not the people who would die in the catastrophe would have existed had another policy been chosen instead. The people who die in the catastrophe resulting from Parfit's Risky Policy thus suffer genuine harms for which we are morally responsible.⁷

And if this is the case, then we can solve the non-identity problem by rejecting P2.

There are several reasons to reject this solution to the problem. The first is that the argument for doing so is guilty of begging the question. Hanser appeals to the claim that the proper account of why we should say that the people in the Same People Choice analogue of the Risky Policy case would be harmed in the morally relevant sense is that they would be harmed by being killed. But this is simply to presuppose that P2' provides a better analysis of the morally relevant sense of harm than does P2. If P2 is correct, then the reason that the choice of the Risky Policy harms people in the morally relevant sense in the Same People Choice variant is simply that it makes them much worse off than they would have been had the policy not been selected. Since on this account, the explanation of why harm occurs in the Same People Choice version does not hold for the choice made in Parfit's original Different People Choice version, we would have no reason to conclude that the people in Parfit's Risky Policy case are harmed in the morally relevant sense. Hanser treats the Risky Policy case as a "counterexample" to P2, but it is a counterexample only if we have an independent reason to believe that the people in the Risky Policy case are harmed in the morally relevant sense. And Hanser provides no such reason.

In addition, there are positive reasons to prefer P2 to P2'. The first reason arises from the familiar fact that everyone who ever lives is harmed in many ways. It follows that any couple who ever considers conceiving a child knows that, if they do conceive, the resulting child will be harmed much more than it would be if they do not conceive (since it would not exist and so would not be harmed at all). This means that P2' entails that every act of conception causes harm in the morally relevant sense. But this is absurd. It may be plausible to maintain that you harm a child by conceiving it if its life will be worse than no life at all, but it is implausible to maintain that you harm it in a morally relevant sense by conceiving it if it will live a typical human life merely because a typical human

being suffers many harms over the course of its life that it would not suffer if it were not conceived. Moreover, the reason this strikes us as implausible is precisely because we recognize that, despite the fact that conceiving a new human being results in its being harmed in various ways that it would not otherwise be harmed, these harms are compensated by the various benefits that it would not otherwise enjoy, such that, on the whole, the child is not made worse off by being conceived. This is why we do not think that conceiving under typical circumstances harms in the morally relevant sense, and this is why P2 is preferable to P2'. A second, and related, reason for preferring P2 to P2' arises from considering cases of life extension. Suppose that a doctor saves a patient's life, knowing that if he does the patient will soon after lose his leg. P2' implausibly entails that the doctor harms the patient in the morally relevant sense by causing him to later lose the leg, which in itself is a harm to the patient. But P2 correctly accounts for the fact that the doctor does not harm the patient in a morally relevant sense by appealing to what seems clearly to be the relevant feature of the case: that although the patient will suffer a harm if he is saved that he will not otherwise suffer, saving him does not make him worse off than he would otherwise be. Finally, if we accepted P2' rather than P2 as an account of causing harm in the morally relevant sense, there would seem to be no clear reason not to accept a parallel account of causing death in the morally relevant sense. We would have to say that you cause a person's death if your act makes him suffer a death rather than not suffer a death. And, since all human beings are mortal, this would implausibly entail that everyone who ever conceives a child kills that child in the morally relevant sense.

2.3. Rejecting P3

P3 stipulates that Wilma's act of conceiving a blind child now rather than a sighted child later does not harm anyone other than the child she conceives. This premise is motivated by the fact that our intuition that Wilma's act is morally wrong does not depend on the act's having negative consequences for any other people. Since we believe that her act would be wrong even if it did not harm anyone else, it will be a problem if we cannot show her act to be wrong without appealing to such secondary harms. It is thus no objection to P3 to point out that, in the real world, conceiving handicapped children often does impose some further costs on other people. The only way to reject P3, then, would be to claim not merely that it is contingently true that some other people are usually harmed by acts such as Wilma's, but rather that it is necessarily true that someone else always is.

I am aware of only two ways in which P3 might be thus attacked, and neither is sufficiently plausible to ground a solution to the non-identity problem. The first possibility is to insist that Wilma's act harms herself. It will be more difficult for her to raise a handicapped child than to raise a non-handicapped child, and so in choosing to conceive now she harms and thus wrongs herself. This possibility must be rejected for two reasons. First, it is not at all clear that it must always

be the case that raising a handicapped child is harmful relative to raising a non-handicapped child. There may be distinctive burdens that the former parent incurs, but there may also be distinctive rewards. But second, and more importantly, even if the choice to conceive a handicapped child does always harm the person who makes the choice, this cannot be what accounts for our intuition that the act of making that choice is wrong. This is because even those who deny that we have any moral duties to ourselves will still agree that Wilma's act is wrong. Of course, we might still think that Wilma's choice is extremely irrational, and we might well disapprove of it (at least in part) on these grounds, but it should be clear upon reflection that not every irrational choice is an immoral choice. Still, in order to avoid being distracted by the seemingly sheer irrationality of Wilma's decision, I will stipulate the following for the purposes of the remainder of this paper: Wilma has two sisters. One has a handicapped child, the other has a non-handicapped child. Wilma has observed both sisters carefully for a long period of time and as a result reasonably believes that the two sisters are equally happy with their children. Neither would trade their child for any other in the world. Having thought about it quite carefully, Wilma is confident that she will be just as happy regardless of whether or not she has a handicapped child. We will presumably still believe that Wilma's act is wrong, but assuming that we do believe this, no appeal to self-regarding duties will be able to account for our belief.

The only other approach that I am aware of here would be to claim that when Wilma chooses to conceive now rather than later, she harms the child she would otherwise have conceived by not conceiving him (since it is a necessary consequence of her act that she not conceive this other child). In choosing to conceive Pebbles, that is, Wilma harms Rocks. But this claim would produce the unacceptable result that every time we decline to reproduce we seriously harm those we could have conceived and that our declining to conceive is thus *prima facie* seriously immoral. These results are far more implausible than the result we attempting to avoid in solving the non-identity problem. And so, again, there proves to be no good reason to reject P3.

2.4. Rejecting P4

P4 maintains that if an act does not harm a person, then it does not wrong that person. Here, at last, we seem to arrive at a genuine weak spot in the argument. For P4 as stated is most likely false. There are many cases in which an act does not harm someone but is at least widely believed to wrong them.

One kind of case involves acts which in fact cause no harm but which are done with the intent to cause harm or at least with extreme indifference to whether or not they cause harm. It is wrong, for example, to point a gun at what you believe to be another person and pull the trigger, even if it turns out that what you shoot is a mannequin. And it would be wrong to fire a gun into a crowded street even if you happened to avoid hitting anyone. Cases such as these suggest that P4 is

false and must be replaced by something like:

P4': If P's act does not harm Q *and* P's act is not done with the intent to harm Q *and* is not done with extreme indifference to whether or not it harms Q, then P's act does not wrong Q

The problem with attempting to solve the non-identity problem by attacking P4 in this way, however, is that we can simply stipulate that Wilma's act has neither of these further wrong-making properties. She does not choose to conceive now in order to cause harm to anyone. She chooses to conceive now in order to avoid the inconvenience of taking the pill once a day for two months. And we can stipulate that she does not act with indifference to whether or not her act will cause harm to anyone. Indeed, we can stipulate that she only decides to conceive now because she has thought things through carefully and has determined that her choice will not harm anyone. Substituting P4' for P4 will therefore do nothing to prevent us from reaching the problematic conclusion that Wilma's act is not wrong.

A second kind of counterexample to P4 involves cases in which an act violates someone's rights but does not harm them. If I break a promise to you, for example, or lie to you, or discriminate against you on the basis of your race, then many people would say that I have treated you wrongly, even if you suffer no harm as a result of my act and even if I knew ahead of time that you would suffer no such harm. It is possible, of course, to respond to such cases by maintaining that a person's life goes less well for him when, for example, he is lied to, so that lying to someone necessarily harms them. But it is important to consider whether the non-identity problem can really be overcome by appeal to such cases even if we concede that they involve acts that wrong people without harming them. So let us assume that there are acts, such as some cases of lies and promise breakings, that wrong people without harming them. The most natural explanation for this would seem to be that some acts are wrong because they violate people's moral rights, and that sometimes they violate people's moral rights without harming them. If this is so, then we must revise P4 even further and replace it with something like this:

P4'': If P's act does not harm Q *and* P's act is not done with the intent to harm Q *and* is not done with extreme indifference to whether or not it harms Q, and if P's act does not violate Q's rights, then P's act does not wrong Q.

Several writers have suggested that this further concession can be used to solve the non-identity problem. Doran Smolkin, for example, argues that cases such as lying and promise breaking demonstrate that there is nothing problematic about saying that people's rights are violated in non-identity cases. When a doctor lies to her dying patient to ease his suffering, for example, this can be wrong because it violates a right that protects the patient's interest in maintaining self-determination and because acting against this particular interest may not be justified even if doing so does not leave the patient worse off in terms of promoting his

interests overall. But if this is so, Smolkin suggests, then “there appears to be nothing necessarily implausible or even particularly unusual about saying that a person can have a legitimate complaint that some act violated her rights even though she may not have been better off had that act not been performed.”⁸ And if this is so, then there is no particular reason to doubt that Wilma’s act might violate someone’s rights even though her act leaves no one worse off than they would otherwise have been.

Even if we accept this further modification of P4, however, there is little reason to believe that this will enable us to solve the non-identity problem. In the case where the doctor lies to her patient, after all, it is easy to identify the person whose right the act violates. And the same is true of the other sorts of cases that writers such as Smolkin appeal to: if I lie to you, or break a promise to you, or discriminate against you on the basis of your race, then it is you whose rights I have violated and it is the right not to be lied to, or to have a promise made to you broken, or to be discriminated against, that I have violated. But even if we agree that acts can be wrong in such cases without harming anyone, it is extremely difficult to see how this could be used to account for the wrongness of Wilma’s act. For it will enable us to account for the wrongness of Wilma’s act only if we can identify some right that someone has that is violated by her act. And it is extremely difficult to imagine what this right could be. We cannot say that Wilma’s act violates Pebbles’s rights, because Pebbles does not exist at the time that her act is performed. We could say that her act violates the rights of the potential Pebbles who does not yet exist at the time of her act, but this is doubly troubling: it is not clear what grounds there are for attributing actual rights to merely potential people, and even if we thought that potential Pebbles could have actual rights, what right would she have that Wilma could be understood to be violating? Since Wilma’s act is the act that results in potential Pebbles being made actual, the only kind of right that Wilma could be understood to be violating would be the right of potential Pebbles not to be conceived into an actual person. But what grounds could there be for saying that potential Pebbles would have such a right? It might be plausible to maintain that this would be so if conceiving Pebbles would give her a life worse than no life at all, but this is hardly the case in conceiving a blind but otherwise healthy child. How could a potential person have a right not to be conceived given that conceiving them would give them a life well worth living?⁹ The only salient alternative would seem to be to say that by conceiving Pebbles, Wilma violates the rights of the potential child she would otherwise have conceived had she taken the pill once a day for two months and then conceived. This would mean that she violates the right of this potential child to be conceived. But it is difficult to imagine how one could hope to ground such a right without it being the case that every child that we have the ability to conceive has a right to be conceived by us. And it is even more difficult to imagine that we could accept such a claim. So even if we replace P4 with P4’ and admit

that there are acts that wrong people by violating their rights without harming them, there seems to be no way to exploit this revision as a way of overcoming the non-identity problem. The problem will simply reappear once we add the further claim that Wilma's act does not violate anyone's rights. And since this premise seems just as secure as the others, the problem seems to remain just as immune to solution.

The attempt to solve the non-identity problem by appealing to the claim that acts that are harmless can nonetheless be morally wrong because they violate rights therefore fails. There is, however, a second and more subtle way in which one might attempt to make use of the notion of rights to overcome the problem. This second approach arises from the idea that if it is wrong to violate a particular right that a particular person has, then it is wrong to generate that right in that person, knowing that the right will be violated.¹⁰ If, for example, it is wrong to violate the right that someone has that you keep a promise that you made to them, then it is wrong for you to generate that right in that person by making a promise to them if you know ahead of time that the promise will not be kept. Similarly, the argument continues, if it is wrong to generate a right in someone that you know will be violated, then it is wrong to create a person with a right that you know will be violated. But in the kind of case in which the non-identity problem arises, the argument concludes, a person does an act of precisely this kind. Wilma's act, on this account, is not wrong because the act itself violates some potential or possible person's rights, but rather because it predictably causes the existence of an actual person with rights that will subsequently be violated. And so, on this account, we can appeal to the importance of moral rights as a way of denying P4 without having to insist that possible or potential people already have actual rights.¹¹

There are several reasons to reject this further rights-based solution to the non-identity problem. These reasons parallel the reasons that were given in section 2.2 to reject Hanser's solution. First, the solution rests on a faulty inference. It assumes that in the case in which you make a promise that you know will not be kept, the reason that your act is wrong is that the person you make the promise to will later have more rights violated than he would otherwise have. Only if this is the reason for the wrongfulness of your act will the reason that your act is wrong carry over to the case in which you create a new person with a right that you know will be violated. But there is a natural alternative explanation for the wrongness of your act of making a promise to someone when you know the promise will not be kept: the act is wrong not because it means that he will later have more rights that are violated than he would otherwise have, but rather because the act itself violates the right that he already has not to be the subject of a false promise. The difference between these two explanations of the wrongness of making a promise you know will not be kept is crucial. The first explanation entails that it is also wrong to create the further person with a right that you know will be violated,

since that act also makes it the case that someone will later have more rights that are violated than they would otherwise have (since they would otherwise not exist and thus have no rights violated). But the second explanation does not have this further implication. On the second explanation, your act is wrong because it violates a right that the person you make the promise to already has, and, as we have already seen, it is implausible to maintain that Wilma's act of conceiving Pebbles is an act that violates a right that someone already has. As in the case of Hanser's argument, then, the wrongness of making a promise you know will not be kept can only be used to undermine the argument that gives rise to the non-identity problem by begging the question. The wrongness of making a promise that we know will not be kept, that is, shows that we wrong people by creating them with rights that will be violated only if we begin by assuming that this is the best explanation of the wrongness of making a promise that we know will not be kept in the first place.

A further problem with this solution to the non-identity problem parallels the further problem that arose for Hanser's solution: it entails that it is morally wrong to conceive children under any set of circumstances. After all, everyone who lives has at least some of their rights violated as they go through life. Everyone is lied to at some point, has a promise to them broken to them at some point, and so on. If we accept the principle underlying this solution, therefore, then it will follow that it is always wrong to conceive new people.¹² And the same will go for cases involving life extension. Finally, even if we accepted the principle offered here, it is still not clear that it could solve the non-identity problem. For it is not at all clear what right it is that Wilma knows that Pebbles will have and that will be violated once she is born. Is it the right to see things? The right to enjoy a certain level of overall well-being? On both the direct and the more indirect approach, then, appealing to the ways in which we can wrong people without harming them fails to solve the problem.¹³

2.5. Rejecting P5

There is one final possibility. We can deny P5. P5 maintains that if an act does not wrong anyone, then the act is not wrong. If we reject P5, then we can admit that Wilma's act does not really wrong anyone but still maintain that her act is morally wrong. Surely we could do this. We could pick out some property other than the property of being an act that harms someone or being an act that wrongs someone and insist that this other property is sufficient to make an act morally wrong. The question, though, is whether we can do this in a way that successfully entails that Wilma's act is morally wrong without also generating further implications that are at least as problematic as the claim that her act is not wrong. And it seems to me that we cannot.

The reason for this is that there seem to be only two further properties to appeal to, and neither provides the basis for an acceptable moral principle. One

further feature of Wilma's act is that it results in the existence of a person with a significant handicap. In order to avoid the conclusion of the non-identity argument, therefore, we could treat the existence of people with significant handicaps as itself an intrinsically bad thing. Gregory Kavka, for example, has proposed that we can avoid the non-identity problem, at least in a good number of cases, by appealing in part to some version of a principle on which "conditions of society or the world are intrinsically undesirable from a moral point of view to the extent that they involve people living restricted lives" where this means living "a life that is significantly deficient in one or more of the major respects that generally make human lives valuable and worth living."¹⁴ We could then reject P5 by saying that an act can be wrong, even if it harms no one and wrongs no one, if it results in the existence of an intrinsically bad thing that would not otherwise exist. Accepting such a principle would indeed permit us to reject P5. But it would do so only by generating consequences that are even more difficult to believe than the one generated by the non-identity argument. We would have to say that every time a blind person is born, for example, the world becomes a worse place. We would have to say that it would be better, from the moral point of view, that such people never exist. And this is clearly unacceptable.

The other potentially salient feature of Wilma's act is that it involves making a choice such that had she made the other choice, the overall resulting state of affairs would have been better. So we could also deny P5 by maintaining that an act can be morally wrong even if it harms no one and wrongs no one if it is done instead of some other possible act that would have resulted in an overall better state of affairs. Neil Levy, for example, in discussing a case virtually identical to that of Wilma, proposes that the problem can be solved by appealing to the claim that "Since her action is not the best of the alternatives available to her, she acts wrongly."¹⁵ But this principle, too, has implications that are even more objectionable than the problem it is invoked to avoid. It would imply, for example, that if the world would be a better place overall if you had a child, then it would be wrong for you to decide not to.¹⁶ If it would be better overall if you were to have a child next year rather than this year, then it would be wrong to have it this year. If it would be better overall for you rent a movie tonight rather than see one in the theater and send the money you save by doing so to Oxfam, then it would be wrong for you to see a movie in a theater tonight. If it would be better for you to donate some money to cancer research than to donate it to diabetes research, then it would be wrong for you to donate the money to diabetes research, and so on.¹⁷

A sufficiently committed consequentialist, of course, might well simply accept such implications. And so a sufficiently committed consequentialist might well be able to avoid the non-identity problem. But avoiding the non-identity problem by embracing consequentialism is not the same thing as solving the non-identity problem. The non-identity problem arises in the first place, after

all, because it reveals an apparent inconsistency within common-sense morality: general moral principles that most people seem to accept turn out to contradict assessments that most people make in such cases as those involving Wilma and the Risky Policy. Abandoning common-sense morality for consequentialism does nothing to eliminate this problem. If anything, by highlighting how drastically consequentialism departs from common-sense moral beliefs, this response highlights just how powerful the non-identity problem is. If we can solve the problem only by accepting the claim that we are obligated to procreate whenever doing so would make the world a better place, after all, then the problem must be powerful indeed. And on the assumption that we are unwilling to embrace such claims, we cannot solve the non-identity problem by rejecting P5.

3. AN ALTERNATIVE APPROACH

I have argued that we cannot avoid the conclusion of the argument that gives rise to the non-identity problem by rejecting one of the argument's premises. If this is so, then what other kind of solution to the problem can we turn to? One possibility would be to reject the conclusion by showing that it does not follow from the premises. But since each of the steps of the argument follows as a matter of elementary logic, there is no way to do this. A second possibility would be to deny the conclusion while accepting the premises and while admitting that logic dictates that the conclusion follows from the premises. This, in effect, seems to be the approach of the British Law Commission's 1974 *Report on Injuries to Unborn Children*, which, in considering cases related to those such as that of Wilma and Pebbles, states, "[We] have not been unduly influenced by these considerations of logic. Law is an artifact and, if social justice requires that there should be a remedy given for a wrong, then logic should not stand in the way."¹⁸ But this, surely, is not to solve the problem but simply to give up in the face of it. A third possibility would be to simply refuse to assess the morality of Wilma's act in the first place by insisting that the question raised by the non-identity problem simply lies, as Heyd suggests, "outside the realm of ethics." Ethical questions, on this analysis, can only concern our relations with other actually existing people, and questions such as that raised by Wilma's decision must therefore be placed instead in some separate "nonmoral sphere of deliberation."¹⁹ But avoiding the non-identity problem by redrawing the lines that distinguish ethical questions from other questions is merely another form of evasion. Regardless of whether we call the question an ethical one, a "genetical" one, or something else, the question of whether or not Wilma is doing something that she ought not to do seems to be a perfectly intelligible one and one that should have a meaningful answer. There seems, then, to be no way to resist the argument's premises, no way to resist the move from the premises to the conclusion, and no justification for simply avoiding the issue.

Is there any other way to solve the problem? There is. For we can simply ac-

cept the argument's conclusion. The reason that the premises outlined in section 1.2 are taken to generate the non-identity *problem* after all, rather than simply the non-identity *argument*, is that we resist the argument's conclusion. Abandon our resistance, and the problem goes away. The non-identity problem, understood as a puzzle that stands in need of resolution, disappears, and is replaced by the non-identity argument, which is simply a sound argument for an admittedly surprising conclusion.

3.1. *The Hypothesis*

Simply accepting the conclusion of the non-identity argument, of course, is easier said than done. We must first identify and confront the source of our resistance to it. So why, precisely, do we resist the claim that Wilma's act is not wrong? It is not because we have a sound independent argument that demonstrates that her act is wrong. Indeed, if anything, the only apparently sound argument we have seems to show that her act is not wrong. Rather, the only reason we have for resisting the claim that Wilma's act is not wrong is the fact that the claim is so counterintuitive. In order to successfully solve the non-identity problem by accepting the argument's conclusion, therefore, we must confront this fact. But I believe that this fact can be successfully confronted.²⁰ For I believe that there is a reason that we find the conclusion of the non-identity argument so counter-intuitive, and that once this reason is understood, we will come to see that we have little reason to endorse our intuitive reaction to the story of Wilma. Rather than providing an accurate reflection of our more basic moral beliefs, that is, our intuitive reaction to cases such as that of Wilma instead produces a distorted view of those beliefs. Once we come to doubt that our intuitive reaction to the Wilma case is a reliable reflection of our more basic moral commitments, we will have no further reason to deny the conclusion of the argument, and the non-identity problem will disappear, leaving in its place simply a sound argument.

My hypothesis about what happens to our intuitions in the Wilma case is as follows: our basic moral beliefs include the belief that it is (prima facie) wrong to act in ways that significantly harm other people. In typical cases, an act that causes a person to have a significant handicap is an act that significantly harms that person. This is because, in typical cases, an act that causes a person to have a significant handicap is an act that causes that person to be significantly worse off than they would otherwise have been. When our moral intuitions register disapproval of an act that involves causing someone to have a significant handicap, this is because the intuition is reflecting this more fundamental moral belief. This is what happens, for example, in the case of Betty.

The case of Wilma, however, is atypical. Because of facts that are idiosyncratic to that case, the act of causing a person to have a significant handicap in the case of Wilma does not, in fact, cause that person (or any person) to be worse off than they would otherwise have been. But the way in which this is so is dif-

difficult for us to keep present before our mind's eye as we contemplate the case. As a result, I suggest, that component of our basic moral values that disapproves of the causing of significant harm to others still registers with us and because of this we still find ourselves feeling that Wilma has done a morally wrong act. At a purely intellectual level, we understand that Wilma's act harms no one, but our moral intuitions are not sufficiently fine tuned to respond at every instance to what, at a purely intellectual level, we understand to be the case. As a result, because the Wilma case so closely resembles other cases that do involve causing significant harm to others, and since the features of the case that account for the fact that Wilma does not cause any harm at all in this instance are difficult to keep prominent in our mind's eye as we reflect on the case, we still feel the force of the intuition that her act is wrong because we feel the force of the general moral belief that causing significant harm to other people is wrong.

Let me try to be a bit more specific about what my hypothesis involves before moving on to discuss how we might test its plausibility. So first, consider the following two statements about Wilma:

1. Wilma's act makes it the case that her child is blind rather than that her child is sighted.
2. Wilma is responsible for the fact that Pebbles is now blind.

Each of these sentences can be interpreted in two different ways. The ambiguity in the first sentence arises from the *de re/de dicto* distinction discussed briefly in section 2.1. The sentence could mean either that (a) there is some specific individual who is blind rather than sighted (the *de re* sense) or that (b) the person picked out by the description "her child" is one individual who is blind rather than a second individual who is or would have been sighted (the *de dicto* sense). The ambiguity in the second sentence arises from the fact that there are two different ways in which it could fail to be the case that Pebbles is now blind: she could exist and be sighted or she could not exist. Thus, the second sentence could mean either that (a) Wilma is responsible for the fact that Pebbles exists blindly rather than exists sightedly, or that (b) Wilma is responsible for the fact that blind Pebbles is existing rather than not existing. In each case, if the sentence is true in the (a) sense, this would mean that Wilma has done something harmful: making a particular child exist as blind who would otherwise exist as sighted. And in each case, if the sentence is true in the (b) sense, this would mean that Wilma had done something harmless: making a child who could only exist in a blind state exist rather than not exist. In each case, upon careful reflection, it should be clear to us that the sentence is, in fact, false in the (a) sense and true in the (b) sense. And so, upon careful reflection, it should be clear to us that Wilma's act has not harmed anyone. But, and this is the crucial point, when we think about these statements, it is difficult for us to keep the (a)/(b) distinctions clear in our mind's eye. It is natural to slip, for example, from thinking that Wilma is responsible for the fact

that Pebbles is blind where this means that had Wilma not done the act Pebbles would not now exist to thinking simply that Wilma is responsible for the fact that Pebbles is blind. And it is natural to slip from thinking this to thinking that it is Wilma's fault that Pebbles is blind rather than sighted.

My hypothesis, then, can be restated more precisely as the view that what accounts for our intuition that Wilma's act is wrong is the fact that even though at an intellectual level we know that statements such as 1 and 2 are true only in the (b) sense, we have difficulty keeping this fact clear in our mind's eye. We are tempted to think that since the sentences are true in the (b) sense, they are, simply, true, and that if they are true, then they must be true in the more straightforward (a) sense as well. The belief that the sentences are true in the (a) sense involves the belief that Wilma's act has caused significant harm to someone, and this belief, in turn, triggers an intuitive reaction to the effect that Wilma's act is wrong. My hypothesis, then, is that it is only in virtue of the fact that we are prone to confuse the (a) and (b) senses of such sentences that we end up experiencing the intuitive judgment that Wilma's act is wrong. At a purely intellectual level, to take a different kind of example, a person awaiting the toss of a fair coin may fully understand that the fact that the previous ten tosses of the coin have all landed on heads does not mean that the coin is "due" to land on tails. But at an intuitive level, it can nonetheless be extremely difficult to avoid feeling that it does. Similarly, we may fully understand at a purely intellectual level that Wilma's act has not harmed anyone. But at an intuitive level, it can be extremely difficult to avoid feeling that she has.

3.2. Testing the Hypothesis

How could we test this hypothesis? My hypothesis can be used to generate a prediction about how we would be likely to respond to a different kind of case. We can therefore devise a test case to see if the prediction is vindicated. If we do not respond to the test case in the way that my hypothesis predicts we will, this will count as evidence against the hypothesis. But if we respond to it in the way that my hypothesis predicts we will, this will count as evidence in its favor. The prediction that arises from my hypothesis is this: suppose we had a case that was just like the Wilma case in that it rendered the two kinds of sentence above true in the (b) sense and false in the (a) sense. But suppose, unlike the Wilma case, the case was such that it was virtually impossible for us to confuse the two senses. Suppose, that is, that in this further case we were not remotely tempted to think that the person had caused any harm by his or her act. In that case, my hypothesis yields a clear prediction: in this further case, we should not have the intuitive reaction that the act is wrong. My hypothesis, that is, maintains that our intuition that Wilma's act is wrong is generated only because we confuse the (a) and (b) senses in which the relevant sentences can be true. Take away the confusion, and we should take away the intuition.

What we need, then, is a further case in which someone makes a choice that is structurally symmetric to the choice that Wilma makes, but in which we will not be at all inclined to be confused into thinking that the choice has harmed anyone. I propose the following:

Fred: Fred is walking past a lake in which two young boys, Billy and Timmy, are drowning. Billy is near Fred while Timmy is further away. Fred has one life preserver. He can toss the life preserver to Billy, but by the time that Billy is safely ashore, it will be too late for Fred to save Timmy. He can run down the shore and toss the life preserver to Timmy, but by the time that Timmy is safely ashore, it will be too late to run back and save Billy. So Fred can only save one of the two boys. Since he would have to walk further to save Timmy, it would be somewhat more convenient for Fred to save Billy. The inconvenience to Fred of saving Timmy rather than Billy is comparable to the inconvenience of taking or giving a tiny pill once a day for two months. Billy is incurably blind, and Fred knows this fact. He also knows that Timmy is not blind. Fred decides to save Billy because this is more convenient for him.

To serve as a fair test case, we must first confirm that Fred's act is morally on a par with Wilma's act. The reason for concluding that this is so is as follows: in both cases, the person has a choice between one of two acts.²¹ In both cases, the choice of which act to do will determine who will exist after the choice has been made. In both cases, if the person makes one choice, an incurably blind child will exist after the choice has been made and a sighted child who would otherwise have existed after the choice has been made will not exist, and in both cases, if the person makes the other choice, a sighted child will exist after the choice has been made and an incurably blind child who would otherwise have existed after the choice has been made will not exist. In both cases, the person chooses to do the act that results in an incurably blind child existing after the choice has been made and a sighted child who would otherwise have existed after the choice was made not existing. And in both cases, the person makes this choice simply because it is somewhat more convenient to do so.²² It is true, of course, that Wilma's choice involves choosing which of two lives to create while Fred's involves choosing which of two lives to extend, but it is difficult to see how this difference in itself could matter morally. If the life of a blind child is valuable enough to be worth saving even when one could instead save a sighted child, after all, how could it not be valuable enough to be worth creating even when one could instead create a sighted child?²³

More specifically, we can confirm that the case of Fred parallels the case of Wilma with respect to the two kinds of sentences discussed above. For here are two sentences about Fred:

1. Fred's act makes it the case that the boy that he saved is blind rather than the boy that he saved is sighted.
2. Fred is responsible for the fact that Billy is now blind.

As with the case of the two sentences about Wilma, each of these sentences can be interpreted in two different ways. The first sentence could mean either that (a) there is a specific boy that Fred's act causes to be blind rather than sighted or that (b) the person picked out by the description "the boy that he saved" is one individual who is blind rather than a second individual who is or would have been sighted. The second sentence could mean either that (a) Fred is responsible for the fact that Billy exists blindly rather than exists sightedly or that (b) Fred is responsible for the fact that blind Billy is existing rather than not existing. And, again as with the case of the two sentences about Wilma, each of the two sentences about Fred is true in the (b) sense and not in the (a) sense.

But, unlike the case of Wilma, it is virtually impossible for us to be confused about the (a)/(b) distinction when we reflect on the case of Fred. The first sentence about Fred is true because "the boy he saved" will refer to one person if Fred saves Billy and to a second person if Fred instead saves Timmy, not because "the boy he saved" will refer to one and the same child regardless of which choice he made, a child who will then be either blind or not blind depending on Fred's choice. And in this case, unlike the case of Wilma, it is inescapably clear to us that the statement is not true in the (a) sense. It is unavoidably obvious that the expression "the boy that he saved" does not refer to the same boy in both instances because both boys already exist in the case of Fred and so we can easily keep them both in front of our mind's eye. We can easily see that if Fred saves the blind boy, then the expression will refer to one child, and that if he saves the sighted boy, then the expression will refer to the other child. But since, in the Wilma case, there is only one child who actually exists (Pebbles), it is not as easy for us to remember that the expression "her child" does not refer to the same child in both the case where she takes the pills and the case where she does not. Indeed, when first presented with the case of Wilma, many people do not even recognize this fact at all. And the same is true of the second sentence about Fred, the statement that Fred is responsible for the fact that the Billy is now blind. We are not the least bit tempted to conflate the (a)/(b) distinction in this case because it is obvious to us that there was nothing that Fred could have done such that Billy would still be alive and not be blind. And, again, this is perfectly clear to us because in the Fred case both boys who could exist after Fred's choice are already present and equally so in our mind's eye. It is just as true in the case of Wilma, however, that there was nothing that Wilma could have done so that Pebbles would now be alive and not be blind. But, once again, this fact is harder to keep prominent in our mind's eye as we contemplate the case.

With respect to both the case of Wilma and the case of Fred, then, both statements are true in the (b) sense and false in the (a) sense. In the case of Fred, it is easy for us to keep both distinctions clear. Indeed, it is virtually impossible to ignore them. As a result, I suggest, the case of Fred provides a good test of my hypothesis. If our response to the story is that what Fred does is morally wrong,

then this will count against my hypothesis. It will show that even when we do not confuse the (a) and (b) senses, we think an act is wrong when it makes such sentences true in the (b) but not (a) sense. And if this is so, then there will be no reason to think that such confusion is what generates our intuition that Wilma's act is wrong. But if our response to the case of Fred is that what Fred does is not at all wrong, then this will count in favor of my hypothesis. It will show that when we are not confused about the (a)/(b) distinction we do not think an act is wrong when it makes such sentences true in the (b) but not (a) sense. And this will strongly suggest that if we were not confused in the case of Wilma, we would not feel that her act was morally wrong.

At this point, of course, I must simply defer to the reader's intuitions. But if my admittedly somewhat limited experience in describing and discussing this case with others is any guide, virtually everyone will agree that what Fred does is not morally wrong. Certainly this is my own reaction to the case of Fred. And if this is so, then this supports my claim that we only respond negatively to Wilma's act because our intuitions get confused by the (a)/(b) distinction. This shows, in turn, that our intuitive response to the Wilma case is unreliable. We have no good reason to endorse it. We may continue to feel some of the force of the intuition that Wilma's act is wrong, of course, just as a straight stick submerged half-way in a glass of water will continue to appear to us to be bent even after we are told why this is so. But just as the explanation of how the water distorts our perception of the stick's shape shows that we should not treat the stick's bent appearance as a reason to think that the stick really is bent, so the explanation of how the facts of Wilma's case distort our perception of her act's wrongness shows that we should not treat our intuition that her act is wrong as a reason to think that her act really is wrong. If this is so, then we have no reason to reject the conclusion of the non-identity argument. And if this is so, then we have no reason to think that there is a non-identity problem.

3.3. A Second Test Case

I have argued that the case of Fred is morally on a par with the case of Wilma, that our intuitive reaction to the case of Fred provides a more reliable indicator of our basic moral values than does our intuitive reaction to the case of Wilma, and that since there seems to be nothing morally wrong with what Fred does, we have no good reason to reject the non-identity argument's conclusion that there is nothing morally wrong with what Wilma does. A critic of my argument, however, might maintain that there is an important difference between the two cases. In particular, he might claim that the difference between life creation and life extension is sufficiently important to undermine the analogy that the argument depends on. I am myself skeptical about whether such a critic can succeed. But it is important to see that my hypothesis can be sustained even if the appeal to the particular example of Fred is ultimately rejected. Sustaining the hypothesis

requires consideration of a further case, one that involves life creation rather than life extension. And although this seems to require the case to be considerably more artificial, in the end it seems to me that the argument itself retains its basic strength. So consider the admittedly somewhat less straightforward case of Barney:

Barney: Barney is walking past a lake in which two young women are drowning. One is near him, the other is further away. Barney has one life preserver. He can toss the life preserver to the woman who is near him, but by the time she is safely ashore, it will be too late to save the further away woman. He can run down the shore and toss the life preserver to the further away woman, but by the time that she is safely ashore, it will be too late to save the first woman. So Barney can only save one of the two women. Since he would have to walk further to save the second woman, it would be somewhat more convenient for Barney to save the woman who is near him. The inconvenience to Barney of saving the further away woman rather than the nearby woman is comparable to the inconvenience of taking or giving a tiny pill once a day for two months. Barney also knows the following facts about the women: each of them has just had sexual intercourse and conception will occur in each of them in approximately an hour, if they do not die first. One of the women will conceive an incurably blind child. The other will conceive a child who is not blind.²⁴ The woman who will shortly conceive a blind child is the one who is closer to Barney. Barney decides to save the woman who will shortly conceive a blind child because this is more convenient for him.

It seems just as clear to me that Barney's act is not morally wrong as it is that Fred's act is not morally wrong. It also seems clear to me that even if one believes that the difference between life creation and life extension undermines the argument based on the Fred example, it cannot undermine a parallel argument based on the Barney example. For the Barney case is just like the Wilma case even in this respect. In both cases, the person has a choice between one of two acts. In both cases, the choice of which act to do will determine who will be conceived after the choice has been made. In both cases, if the person makes one choice, an incurably blind child will be conceived after the choice has been made and a sighted child who would otherwise be conceived will not be conceived, and in both cases, if the person makes the other choice, a sighted child will be conceived after the choice has been made and an incurably blind child who would otherwise be conceived will not be conceived. In both cases, the person chooses to do the act that results in an incurably blind child being conceived after the choice has been made and a sighted child not being conceived who otherwise would have been conceived. And in both cases, the person makes the choice simply because it is more convenient for them to do so. So it seems clear to me that Wilma's choice is morally wrong if and only if Barney's is morally wrong.

Let us now suppose that I am right about this. The Wilma case produces one intuition, the Barney case produces a very different intuition, but our judgment

about the two cases should really be the same. Which intuition, then, should we accept? Which case does a more reliable job of capturing our deepest moral convictions? It seems just as clear here as in the case of Fred that Wilma's case involves a distorting feature that Barney's does not. In the case of Wilma, as I have already suggested, our intuitions are deceived by the feeling that Wilma's act has somehow made her blind child worse off than she would otherwise have been. Despite the fact that on an intellectual level we know this to be false, it is still natural for us to react to her case as if it were true. It is true, for example, that either "her child" will be blind or "her child" will not be blind, and although we understand that the expression "her child" does not refer to the same individual in each outcome, it is difficult to focus our moral sentiments to reflect this fact.

But it is easy to focus our moral intuitions to reflect this fact in the case of Barney. In the case of Barney, just as in the case of Wilma, there is a choice to be made that will determine which conception will occur, the conception of a blind child or the conception of a sighted child. But since, in the case of Barney, there are two different women, one of whom will conceive a blind child and the other of whom will conceive a sighted child, depending on the choice that Barney makes, it is inescapably clear and obvious to us that Barney's choice to bring about the conception of the blind child in no way harms that child. As a result, the distorting feature that is present in the Wilma case is absent in the Barney case and we should use our reaction to the Barney case to inform our considered judgment about the Wilma case, rather than the other way around. And since the Barney case, unlike the Fred case, involves a choice about life creation rather than life extension, the argument for this conclusion can be sustained even if one has qualms about appealing to the case of Fred.

4. CONCLUSION

The non-identity problem is a problem only so long as we think we have good reason to resist the conclusion of the argument that gives rise to it. The only reason we think we have to resist the conclusion is that we have strong moral intuitions that run contrary to it. But these moral intuitions can be explained away. They arise not because they accurately reflect our more basic beliefs about the behavior in question, but rather because our moral sentiments are such that they produce misleading results in such cases. This claim is supported by the fact that our intuitions are different in other cases which are otherwise morally similar but which lack the distorting features of the cases that give rise to the apparent problem. Once this becomes clear, we have no reason to reject the conclusion of the non-identity argument. The non-identity problem is solved not by finding a mistake in the argument that generates it, but rather by finding that there is no reason to treat the argument as giving rise to a problem in the first place.

NOTES

1. See Derek Parfit, "On Doing the Best for Our Children," in *Ethics and Population*, ed. Michael Bayles (Cambridge, Mass.: Schenkman, 1976), pp. 100–102; Thomas Schwartz, "Obligations To Posterity," in *Obligations to Future Generations*, ed. Richard Sikora and Brian Barry (Philadelphia: Temple University Press, 1978), pp. 3–13; Robert M. Adams, "Existence, Self-Interest, and the Problem of Evil," *Nous*, vol. 13; and Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), chap. 16.

2. Strictly speaking, there is one further difference: Betty is deterred by the inconvenience of giving a pill to her baby while Wilma is deterred by the inconvenience of taking a pill herself. But I will just stipulate that there is no difference in the level of inconvenience between the two, and since the pill has no side effects it does not matter that Betty would be giving it to someone else rather than to herself.

3. Matthew Hanser, "Harming Future People," *Philosophy and Public Affairs*, vol. 19, no. 1 (Winter 1990), pp. 47–70.

4. Hanser puts it in terms of the distinction between "A choice does not harm someone in the morally relevant sense unless it makes him worse off than he would have been had the choice not been made" and "An act does not harm someone in the morally relevant sense unless its consequence is that that person is harmed more than he would have been had the act not been performed" ("Harming Future People," p. 55).

5. Parfit, *Reasons and Persons*, pp. 371–372; Hanser, "Harming Future People," p. 56.

6. Toward the end of the article, Hanser concedes that even if his solution accounts for the wrongness of the act of choosing the Risky Policy, it will not account for the wrongness of acts like that of Wilma's. He maintains, however, that this is because what is still objectionable about Wilma's choice involves something other than the non-identity problem ("Harming Future People," p. 69).

7. Hanser, "Harming Future People," p. 58.

8. Doran Smolkin, "The Non-Identity Problem and the Appeal to Future People's Rights," *The Southern Journal of Philosophy*, vol. 32 (1994), p. 319. James Woodward offers a similar argument in "The Non-Identity Problem" (*Ethics*, vol. 96 [July 1986], pp. 804–831), and appeals to examples involving promise breaking and racial discrimination. Nils Holtug also suggests that the problem can be avoided by appealing to "deontological principles, e.g., to the effect that certain acts are wrong, even if they harm no-one," though he does not develop this suggestion ("Good for Whom?" *Theoria*, vol. 69, nos. 1–2 [2003], p. 7).

9. In "Toward A Rights-Based Solution to the Non-Identity Problem," Smolkin suggests a possible content for the right that Pebbles might have: the right not to be brought into existence if "a necessary condition for her coming into being also resulted in her being unable to acquire to a sufficient degree one or more of the elements needed for well-being in a particular stage of life" (*Journal of Social Philosophy*, vol. 30, no. 1 [Spring 1999], p. 206). It does seem clear that if we thought that potential Pebbles had this right, then we could say that Wilma's act of conceiving her violated this right. But Smolkin provides no reason to think that potential people have this (or any other right) and it is difficult to see what reason there could be. We could, of course, simply stipulate that this right exists simply in order to solve the non-identity problem. But we could just as easily stipulate something else to render false one of the other premises of the argument that generates the

problem. Without a reason for believing that potential people have actual rights and that among their rights are rights such as the one that Smolkin suggests, we have no reason to accept this rights-based solution to the problem.

10. This argument is defended by Woodward, "Non-Identity Problem," pp. 812–813ff. See also Michael Tooley, *Abortion and Infanticide* (Oxford: Clarendon Press, 1983), pp. 270–273.

11. Woodward, for example, writes that "I agree that it is a necessary condition for a right or an obligation to be violated that someone holding that right or obligation actually exist. But it is consistent with this claim to hold that the reason a certain course of action would be wrong is that it would involve the creation of rights and obligations that would probably or inevitably be violated, and *this* is what I have claimed with regard to the Risky Policy" ("Non-Identity Problem," p. 821).

12. This problem is noted by David Heyd, *Genethics: Moral Issues in the Creation of People* (Berkeley: University of California Press), pp. 243–244, n. 16.

13. One still further means of denying P4 also mentions brief notice here: that of the contractarian moral theorist. Rivka M. Weinberg, for example, has defended a view on which children such as Pebbles can complain that "their parents *wronged* them by not abiding by just procreative principles" where such principles are understood to be those that would be agreed to under suitably fair and impartial conditions, and that this can be so even if their parents' acts do not leave them worse off than they would otherwise have been ("Procreative Justice: A Contractualist Account," *Public Affairs Quarterly*, vol. 16, no. 4 [October 2002], p. 407). On Weinberg's account, the parties to the contract that agree to principles of procreative justice are those people who would exist if the principles under consideration were adopted and enforced (i.e., they would be conceived in accordance with the principles and would conceive in accordance to the principle) and they reach this agreement under a suitably specific veil of ignorance (pp. 409–410).

This contractarian solution can be handled quickly because it has already been satisfactorily rebutted by Parfit (see also Heyd, *Genethics*, pp. 41–51). Parfit points out that the application of contractarian reasoning to decisions involving procreation leads to results that are even more problematic than the judgments produced by the non-identity problem. In choosing along the lines that Weinberg suggests, for example, we would end up preferring what Parfit calls Hell I, in which a great many people suffer for one hour less than fifty days, over what he calls Hell II, in which considerably fewer people suffer for a full fifty days. If I knew that I would exist in whichever Hell was chosen, I would prefer to exist in the Hell in which the suffering was a bit shorter, but it is clearly absurd to think that this world is the better of the two.

Weinberg attempts to circumvent this decisive objection by complaining that "Parfit is not applying the Rawlsian method correctly. Contractualism is directed toward principles, not outcomes" (p. 409). But the basis of Parfit's rejoinder is just as strong even if we redirect it toward choices involving principles. Suppose, for example, that we were choosing between two principles that would govern our choices in preventing certain kinds of conception from occurring, those which would result in lives containing a great deal of suffering. Principle I would direct us to try to minimize the amount of suffering that any one future person will have, while Principle II would direct us to try to minimize the total amount of suffering that future people would have. Principle I, for example, would deem it more important to prevent the conception of one child who would suffer for fifty days and then die than to prevent the conception of triplets each of whom would suffer for one

hour less than fifty days and then die, while Principle II would deem it more important to prevent the conception of the triplets. Even if we follow Weinberg's edict that contractarian reasoning be applied to the choice of principles rather than outcomes, we will conclude that we should choose Principle I over Principle II (if we know we will exist in either outcome we would prefer to suffer one hour less rather than one hour more) and this is no less problematic than choosing Hell I over Hell II in Parfit's original argument.

14. Gregory S. Kavka, "The Paradox of Future Individuals," *Philosophy and Public Affairs*, vol. 11, no. 2 (1981), p. 105. Dan W. Brock tentatively puts forth a similar proposal in terms of the notion of "limited opportunity" ("The Non-Identity Problem and Genetic Harms: The Case of Wrongful Handicaps," *Bioethics*, vol. 9, nos. 3–4 [1995], pp. 272–275). Kavka also appeals to a modified version of the categorical imperative that "would forbid treating rational beings *or their creation* (that is, their being brought into existence) as a means only, rather than as ends in themselves" as a means of rejecting P5 (p. 110). But while Kavka characterizes this second principle as "independent" of the first, it in fact works to solve the non-identity problem only if the first principle does. The modified version of the categorical imperative is designed to block "the attempt to justify (or excuse) acts whose immorality would otherwise be evident *because* they lead to the creation of restricted lives, on the grounds that those created are benefited (or not harmed)" (pp. 110–111, emphasis added). It can therefore succeed only if the fact that an act involves the creation of a restricted life is, indeed, a sufficient reason to deem the act immoral, and this is what the first principle maintains. If, as I suggest above, Kavka's first principle is unsuccessful, therefore, so too is the second.

15. Neil Levy, "The Apology Paradox and the Non-Identity Problem," *The Philosophical Quarterly*, vol. 52, no. 208 (July 2002), p. 366. In Levy's example, the woman must take the medication for three months before conceiving rather than two, and she seems to decline to take the medicine for no reason at all, rather than to avoid a minor cost to herself, but aside from this the case is essentially the same as that of Wilma. Levy's solution to the problem involves an appeal to temporal indexing in addition to the principle cited in the text, but the principle itself is needed in order for Levy's solution to succeed, and so if my argument against the principle is satisfactory, it is not necessary to address the question of temporal indexing independently.

16. This problem is noted by, e.g., Kavka ("Paradox," pp. 99, 104).

17. It is worth noting in this context that the closest Parfit himself comes to solving the problem is with his proposed principle Q: "If in either of two outcomes the same number of people would ever live, it would be bad if those who live are worse off, or have a lower quality of life, than those who would have lived" (*Reasons and Persons*, p. 361). But even if we accept Q, it cannot solve the problem. It could solve the problem if the problem were that the argument seemed to prove that Wilma does not make the *worse* choice. Q could certainly explain why Wilma's choice is the worse one. But, as I argued in section 1.3, the real problem is that the argument seems to prove that Wilma does not make a morally *wrong* choice, and principle Q does nothing to conflict with this contention. Q could only be used to solve the problem if it insisted that an act was positively immoral if it resulted in there being less overall good in the world and, as noted above, that claim is surely too strong to be acceptable.

18. Quoted by Heyd, *Genethics*, p. 34.

19. Heyd, *Genethics*, p. 66. Heyd characterizes the kind of question raised by Wilma's act as lying outside the domain of ethics throughout his book (e.g., pp. 23, 25, 64, 70, 76, 82, 88, 126).

20. So far as I am aware, no one has attempted to do this. Schwartz, it is true, endorses the conclusion of the non-identity argument in "Obligations to Posterity," but he does nothing to respond to the objection that the argument's conclusion is objectionably counterintuitive. Heyd argues that a variety of empirical facts "can remove the sting of repugnance from many of the hypothetical counterintuitive conclusions" that arise if one concedes all the premises of the argument (*Genethics*, p. 195, see also pp. 15, 79, 209), but this shows only that the non-identity problem can generally be avoided in practice since P3 will generally be false. It does nothing to help us to accept the conclusion if P3 is assumed to be true.

21. I am assuming that Wilma has already decided that she will conceive either now or later and that Fred has already decided that he will save one child or the other.

22. It might be objected that I am overlooking an important difference between the cases of Fred and Wilma. In the case of Wilma, she will be the mother of the child who exists after her choice is made whereas in the case of Fred, he will be unrelated to the child who exists after his choice is made. Since it is commonly believed that we have stronger duties to our own children than we have to other children, it might seem that Wilma has an obligation to ensure that the child she chooses be as healthy as possible, whereas Fred has no such obligation. But there are two problems with this objection. The first is that it simply begs the question about the content of parental duties. While it is generally accepted that we have stronger duties of care to our already-existing children, this does not, in itself, show that we have duties to bring about the existence of healthier children. The second problem is that we can revise the case of Fred to make it just like the case of Wilma with respect to relatedness, and it seems to be just as clear that Fred's act would not be wrong. Suppose, that is, that the two children drowning in the lake are his own children, that he can only save one of them, and that it is easier for him to save the blind child. I doubt that anyone would claim that Fred is obligated to save his sighted child rather than his blind child, but unless he is so obligated, his act is not wrong even in this modified version of the story and so the difference between the case of Fred and Wilma proves to be morally irrelevant.

23. In saying this, I do not mean to insist that there are no morally relevant differences between life creation and life extension. It seems plausible, for example, to suppose that people have a positive duty to save lives that are already being lived but no such duty to bring new lives into existence. But this difference does not affect the argument made here. We can assume that Fred is obligated to save one child or the other (i.e., that it would be wrong for him to allow both to die) while Wilma is not obligated to conceive one child or the other (i.e., that it would be perfectly permissible for her to decline to conceive at all). This does not affect the use I make of the cases here because Fred and Wilma have each already decided that they are going to save/conceive one child or the other and are merely deciding which it will be. And even if it is true that Fred must save one child while Wilma need not conceive any child at all, this can provide no support for the claim that Wilma's act is wrong while Fred's is not. Indeed, if anything, this asymmetry makes it all the more difficult to maintain this view. If Fred is obligated to save a child, after all, then presumably this is because Fred has an obligation to prevent harm from occurring, at least if the prospective harm is significant and the cost to him of preventing it is very

small. But if this is so, then since we are assuming that death will deprive the sighted child of a better life than death will deprive the blind child of, the fact that Fred is obligated to save a child might be used to show that he must save the sighted child. But since Wilma is perfectly free to conceive no child at all, this would make it seem even harder to justify the claim that it is wrong for her to conceive Pebbles. If she is free to introduce no new life into the world, after all, then how can the fact that she could have conceived an even happier child than Pebbles make it wrong for her to conceive Pebbles?

24. Since our intuitions about this case might possibly be affected by considerations involving desert, I will also stipulate that the woman who is about to conceive an incurably blind child had no reason to think that this would be the case when she engaged in sexual intercourse. Even if we assume that intentionally (or negligently) conceiving a blind child is morally wrong, then, the two women drowning in the lake are equal in terms of moral desert.