

Energy Policy and the Further Future: The Identity Problem

DEREK PARFIT

I have assumed that our acts may have good or bad effects in the further future.¹ Let us now examine this assumption. Consider first

The Nuclear Technician: Some technician lazily chooses not to check some tank in which nuclear wastes are buried. As a result there is a catastrophe two centuries later. Leaked radiation kills and injures thousands of people.

We can plausibly assume that, whether or not this technician checks this tank, the same particular people would be born during the next two centuries. If he had chosen to check the tank, these same people would have later lived, and escaped the catastrophe.

Is it morally relevant that the people whom this technician harms do not yet exist when he makes his choice? I have assumed here that it is not. If we know that some choice either may or will harm future people, this is an objection to this choice even if the people harmed do not yet exist. (I am to blame if I leave a man-trap on my land, which ten years later maims a five-year-old child.)

Consider next

The Risky Policy: Suppose that, as a community, we have a choice between two energy policies. Both would be completely safe for at least two centuries, but one would have certain risks for the further future. If we choose the Risky Policy, the standard of living would be somewhat higher over the next two centuries. We do choose this policy. As a result there is a similar catastrophe two centuries later, which kills and injures thousands of people.

Unlike the Nuclear Technician's choice, our choice between these policies affects who will be later born. This is not obvious, but is on reflection clear.

Our identity in fact depends partly on when we are conceived. This is so on both the main views about this subject. Consider some particular person, such as yourself. You are the n th child of your mother, and you were conceived at time t . According to one view, you could not have grown from a different pair of cells. If your mother had conceived her n th child some months earlier or later, that child would *in fact* have grown from a different pair of cells, and so would not have been you.

According to the other main view, you could have grown from different cells, or even had different parents. This would have happened if your actual parents had not conceived a child when they in fact conceived you, and some other couple had conceived an extra child who was sufficiently *like* you, or whose life turned out to be sufficiently like yours. On this other view, that child would have been you. (Suppose that Plato's actual parents never had children, and that some other ancient Greek couple had a child who wrote *The Republic*, *The Last Days of Socrates*, and so on. On this other view, this child would have been Plato.) Those who take this other view, while believing that you *could* have grown from a different pair of cells, would admit that this would not *in fact* have happened. On both views, it is in fact true that, if your mother had conceived her n th child in a different month, that child would not have been you, and you would never have existed.

It may help to shift to this example. A fourteen-year-old girl decides to have a child. We try to change her mind. We first try to persuade her that, if she has a child now, that will be worse for her. She says that, even if it will be, that is her affair. We then claim that, if she has a child now, that will be worse for her child. If she waits until she is grown up, she will be a better mother, and will be able to give her child a better start in life.

Suppose that this fourteen-year-old rejects our advice. She has a child now, and gives him a poor start in life. Was our claim correct? Would it have been better for him if she had taken our advice? If she had, he would never have been born. So her decision was worse for him only if it is against his interests to have been born. Even if this makes sense, it would be true only if his life was so wretched as to be worse than nothing. Assume that this is not so. We must then admit that our claim was false. We may still believe that this girl should have waited. That would have been better for her, and the different child she would have had later would have received a better start in life. But we cannot claim that, in having *this* child, what she did was worse for *him*.

Return now to the choice between our two energy policies. If we

choose the Risky Policy, the standard of living will be slightly higher over the next two centuries. This effect implies another. It is not true that, whichever policy we choose, the same particular people will exist two centuries later. Given the effects of two such policies on the details of our lives, it would increasingly over time be true that people married different people. More simply, even in the same marriages, the children would increasingly be conceived at different times. (Thus the British Miners' Strike of 1974, which caused television to close down an hour early, thereby affected the timing of thousands of conceptions.) As we have seen, children conceived at different times would in fact be different children. So the proportion of those later born who would owe their existence to our choice would, like ripples in a pool, steadily grow. We can plausibly assume that, after two centuries, there would no one living who would have been born whichever policy we chose. (It may help to think of this example: how many of us could truly claim, "Even if railways had never been invented, I would still have been born?")

In my imagined case, we choose the Risky Policy. As a result, two centuries later, thousands of people are killed and injured. But if we had chosen the alternative Safe Policy, these particular people would never have existed. Different people would have existed in their place. Is our choice of the Risky Policy worse for anyone?

We can first ask, "Could a life be so bad—so diseased and deprived—that it would not be worth living? Could a life be even worse than this? Could it be worse than nothing, or as we might say "worth not living?" We need not answer this question. We can suppose that, whether or not lives could be worth not living, this would not be true of the lives of the people killed in the catastrophe. These people's lives would be well worth living. And we can suppose the same of those who mourn for those killed, and those whom the catastrophe disables. (Perhaps, for some of those who suffer most, the rest of their lives would be worth not living. But this would not be true of their lives as a whole.)

We can next ask: "If we cause someone to exist, who will have a life worth living, do we thereby benefit this person?" This is a difficult question. Call it the question whether *causing to exist can benefit*. Since the question is so difficult, I shall discuss the implications of both answers.

Because we chose the Risky Policy, thousands of people are later killed or injured or bereaved. But if we had chosen the Safe Policy these particular people would never have existed. Suppose we do not believe that causing to exist can benefit. We should ask, "If particular people live lives that are on the whole well worth living, even though they are struck by some catastrophe, is this worse for these people than if they had never existed?" Our answer must be "no." If we believe that causing to exist can benefit, we can say more. Since the people struck by the catastrophe live lives that are well

worth living and would never have existed if we had chosen the Safe Policy, our choice of the Risky Policy is not only not worse for these people, it benefits them.

Let us now compare our two examples. The Nuclear Technician chooses not to check some tank. We choose the Risky Policy. Both these choices predictably cause catastrophes, which harm thousands of people. These predictable effects both seem bad, providing at least some moral objection to these choices. In the case of the technician, the objection is obvious. His choice is worse for the people who are later harmed. But this is not true of our choice of the Risky Policy. Moreover, when we understand this case, we know that this is not true. We know that, even though our choice may cause such a catastrophe, it will not be worse for anyone who ever lives.

Does this make a moral difference? There are three views. It might make all the difference, or some difference, or no difference. There might be no objection to our choice, or some objection, or the objection may be just as strong.

Some claim

Wrongs Require Victims: Our choice cannot be wrong if we know that it will be worse for no one.

This claim implies that there is no objection to our choice. We may find it hard to deny this claim, or to accept this implication.

I deny that wrongs require victims. If we know that we may cause such a catastrophe, I am sure that there is at least some moral objection to our choice. I am inclined to believe that the objection is just as strong as it would have been if, as in the case of the Nuclear Technician, our choice would be worse for future people. If this is so, it is morally irrelevant that our choice will be worse for no one. This may have important theoretical implications.

Before we pursue the question, it will help to introduce two more examples. We must continue to assume that some people can be worse off than others, in morally significant ways, and by more or less. But we need not assume that these comparisons could be even in principle precise. There may be only rough or partial comparability. By "worse off" we need not mean "less happy." We could be thinking, more narrowly, of the standard of living, or, more broadly, of the quality of life. Since it is the vaguer, I shall use the phrase "the quality of life." And I shall extend the ordinary use of the phrase "worth living." If one of two groups of people would have a lower quality of life, I shall call their lives to this extent "less worth living."

Here is another example:

Depletion: Suppose that, as a community, we must choose whether to deplete or conserve certain kinds of resources. If we choose Depletion, the quality of life over the next two centuries would be




Figure 10.1. Effects of Choice on Future Standard of Living

slightly higher than it would have been if we had chosen Conservation, but it may later be much lower. Life at this much lower level would, however, still be well worth living. The effects might be shown as in Figure 10.1.

This case raises the same problem. If we choose Depletion rather than Conservation, this will lower the quality of life more than two centuries from now. But the particular people who will then be living would never have existed if instead we had chosen Conservation. So our choice of Depletion is not worse for any of these people. But our choice will cause these people to be worse off than the different people who, if we had chosen Conservation, would have later lived. This seems a bad effect, and an objection to our choice, even though it will be worse for no one.

Would the effect be worse, having greater moral weight, if it was worse for people? One test of our intuitions may be this. We may remember a time when we were concerned about effects on future generations, but had overlooked my point about personal identity. We may have thought that a policy like Depletion would be against the interests of future people. When we saw that this was false, did we become less concerned about effects on future generations?

I myself did not. But it may help to introduce a different example. Suppose there are two rare conditions X and Y, which cannot be detected without special tests. If a pregnant woman has condition X, this will give to the child she is carrying a certain handicap. A simple treatment would prevent this effect. If a woman has condition Y when she becomes pregnant, this will give to the child she conceives the same particular handicap. Condition Y cannot be treated, but always disappears within two months. Suppose next that we have planned two medical programs, but there are funds for only one; so one must be canceled. In the first program, millions of women would be tested during pregnancy. Those found to have condition X would be treated. In the second program, millions of women would be tested when they intend to try to become pregnant. Those found to have condition Y would be warned to postpone conception for at least two months. We are able to predict that these two programs would achieve results in as many cases. If there is

Pregnancy Testing, 1,000 children a year would be born normal rather than handicapped. If there is Pre-Conception Testing, there would each year be born 1,000 normal children, rather than 1,000 different handicapped children. Would these two programs be equally worthwhile?

Let us note carefully what the difference is. As a result of either program, 1,000 couples a year would have a normal rather than a handicapped child. These would be different couples, on the two programs. But since the numbers would be the same, the effects on parents and on other people would be morally equivalent. The only difference lies in the effects on the children. Note next that, in judging these effects, we need have no view about the moral status of a fetus. We can suppose that it would take a year before either kind of testing could begin. When we choose between the two programs, none of the children has yet been conceived. And all of the children will become adults. So we are considering effects, not on present fetuses, but on future people. Assume next that the handicap in question, though it is not trivial, is not so severe as to make life doubtfully worth living. Even if it can be against our interests to have been born, this would not be true of those born with this handicap.

Since we cannot afford both programs, which should we cancel? Under one description, both would have the same effects. Suppose that conditions X and Y are the only causes of this handicap. The incidence is now 2,000 a year. Either program would halve the incidence; the rate would drop to 1,000 a year. The difference is this. If we decide to cancel Pregnancy Testing, those who are later born handicapped would be able to claim, "But for your decision, I would have been normal." Our decision will be worse for all these people. If instead we decide to cancel Pre-Conception Testing, there will later be just as many people who are born with this handicap. But none of these could truly claim, "But for your decision, I would have been normal." But for our decision, they would never have existed; their parents would have later had different children. Since their lives, though handicapped, are still worth living, our decision will not be worse for any of these people.

Does this make a moral difference? Or are the two programs equally worthwhile? Is all that matters morally how many future lives will be normal rather than handicapped? Or does it also matter whether these lives would be lived by the very same people?

I am inclined to judge these programs equally worthwhile. If Pre-Conception Testing would achieve results in a few more cases, I would judge it the better program. This matches my reactions to the questions asked above about our choice of the Risky Policy or of Depletion. There too, I think it would be bad if there would later be a catastrophe, killing and injuring thousands of people, and bad if there would later be a lower quality of life. And I think that it would

not be worse if the people who later live would themselves have existed if we had chosen the Safe Policy or Conservation. The bad effects would not be worse if they had been, in this way, worse for any particular people.

Let us review the argument so far. If we choose the Risky Policy or Depletion, this may later cause a predictable catastrophe, or a decline in the quality of life. We naturally assume that these would be bad effects, which provide some objection to these two choices. Many think the objection is that our choices will be worse for future people. We have seen that this is false. But does this make a moral difference? There are three possible answers. It might make all the difference, or some difference, or no difference at all. When we see that our choice will be worse for no one, we may decide that there is no objection to this choice, or that there is less objection, or that the objection is just as strong.

I incline to the third answer. And I give this answer in the case of the medical programs. But I know some people who do not share my intuitions. How can we resolve this disagreement? Is there some familiar principle to which we can appeal?

Return to the choice of the Risky Policy, which may cause a catastrophe, harming thousands of people. It may seem irrelevant here that our choice will not be worse for these future people. Can we not deserve blame for causing harm to others, even when our act is not worse for them? Suppose that I choose to drive when drunk, and in the resulting crash cause you to lose a leg. One year later, war breaks out. If you had not lost this leg, you would have been conscripted, and been killed. So my drunken driving saves your life. But I am still morally to blame.

This case reminds us that, in assigning blame, we must consider not actual but predictable effects. I knew that my drunken driving might injure others, but I could not know that it would in fact save your life. This distinction might apply to the choice between our two policies. We know that our choice of the Risky Policy may impose harm on future people. Suppose next that we have overlooked the point about personal identity. We mistakenly believe that, whichever policy we choose, the very same people will later live. We may therefore believe that, if we choose the Risky Policy, this may be worse for future people. If we believe this, our choice can be criticized. We can deserve blame for doing what we *believe* may be worse for others. This criticism stands even if our belief is false—just as I am as much to blame even if my drunken driving will in fact save your life.

Now suppose, however, that we have seen the point about personal identity. We realize that, if we choose the Risky Policy, our choice will not be worse for those people whom it later harms. Note that this is not a lucky guess. It is not like predicting that, if I cause you to lose a leg, that will later save you from death in the trenches.

We know that, if we choose the Risky Policy, this may impose harms on several future people. But we also know that, if we had chosen the Safe Policy, those particular people would never have been born. Since their lives will be worth living we know that our choice will not be worse for them.

If we know this, we cannot be compared to a drunken driver. So how should we be criticized? Can we deserve blame for causing others to be harmed, even when we know that our act will not be worse for them? Suppose we know that the harm we cause will be fully compensated by some benefit. For us to be sure of this, the benefit must clearly outweigh the harm. Consider a surgeon who saves you from blindness, at the cost of giving you a facial scar. In scarring you, this surgeon does you harm. But he knows that his act is not worse for you. Is this enough to justify his decision? Not quite. He must not be infringing your autonomy. But this does not require that you give consent. Suppose that you are unconscious, so that he is forced to choose without consulting you. If he decides to operate, he would here deserve no blame. Though he scars your face, his act is justified. It is enough for him to know that his act will not be worse for you.

If we choose the Risky Policy, this may cause harm to many people. Since these will be future people, whom we cannot now consult, we are not infringing their autonomy. And we know that our choice will not be worse for them. Have we shown that, in the same way, the objection has been met?

The case of the surgeon shows only that the objection might be met. The choice of the Risky Policy has two special features. Why is the surgeon's act not worse for you? Because it gives you a compensating benefit. Though he scars your face, he saves you from going blind. Why is our choice of the Risky Policy not worse for those future people? Because they will owe their existence to this choice. Is this a compensating benefit? This is a difficult question. But suppose that we answer "no." Suppose we believe that to receive life, even a life worth living, is not to be benefited.² There is then a special reason why, if we choose the Risky Policy, this will not be worse for the people who will later live.

Here is the second special feature. If we had chosen the Safe Policy, different people would have later lived. Let us first set aside this feature. Let us consider only the people who, given our actual choice, will in fact later live. These will be the only actual people whom our choice affects. Should the objection to our choice appeal to the effects on these people? Because of our choice, they will later suffer certain harms. This seems to provide an objection. But they owe their existence to this same choice. Does this remove the objection?

Consider a second case involving a fourteen-year-old girl. If this second girl has a child now, she will give him a poor start in life. But

suppose she knows that, because she has some illness, she will become sterile within the next year. Unless she has a child now, she can never have a child. Suppose that this girl chooses to have a child. Can she be criticized? She gives her child a poor start in life. But she could not have given *him* a better start in life, and his life will still be worth living. The effects on him do not seem to provide an objection. Suppose that she could also reasonably assume that, if she has this child, this would not be worse for other people. It would then seem that there is no objection to this girl's choice—not even one that is overridden by her right to have a child.

Now return to our earlier case of a fourteen-year-old girl. Like the second girl, the first girl knows that, if she has a child now, she will give him a poor start in life. But she could wait for several years and have another child, who would have a better start in life. She decides not to wait, and has a child now. If we consider the effects only on her actual child, they are just like those of the second girl's choice. But the first girl's choice surely can be criticized. The two choices differ, not in their effects on the actual children, but in the alternatives. How could the second girl avoid having a child to whom she would give a poor start in life? Only by never having a child. That is why her choice seemed not to be open to criticism. She could reasonably assume that her choice would not be worse either for her actual child or for other people. In her case, that seems all we need to know. The first girl's choice has the same effects on her actual child, and on others. But *this* girl could have waited, and given some later child a better start in life. This is the objection to her choice. Her actual child is worse off than some later child would have been.

Return now to the choice between our two social policies. Suppose that we have chosen the Risky Policy. As a result, those who later live suffer certain harms. Is this enough to make our choice open to criticism? I suggest not. Those who later live are like the actual children of the two girls. They owe their existence to our choice, so its effects are not worse for them. The objection must appeal to the alternative.

This restores the second feature that we set aside above. When we chose the Risky Policy, we imposed certain harms on our remote descendants. Were we like the second girl, whose only alternative was to have no descendants? If so, we could not be criticized. But this is not the right comparison. In choosing the Risky Policy, we were like the first girl. If we had chosen the Safe Policy, we would have had different descendants, who would not have suffered such harms.

The objection to our choice cannot appeal only to effects on those people who will later live. It must mention possible effects on the people who, if we had chosen otherwise, would have later lived. The objection must appeal to a claim like this:

(A) It is bad if those who live are worse off than those who might have lived.

We must claim that this is bad even though it will be worse for no one.

(A) is not a familiar principle. So we have not solved the problem that we reached above. Let us remember what that was. If we choose the Risky Policy, or Depletion, this may later cause a catastrophe, or a decline in the quality of life. These seemed bad effects. Many writers claim that, in causing such effects, we would be acting against the interests of future people. Given the point about personal identity, this is not true. But I was inclined to think that this made no moral difference. The objection to these two choices seemed to me just as strong. Several people do not share my intuitions. Some believe that the objections must be weaker. Others believe that they disappear. On their view, our choice cannot be morally criticized if we know that it will be worse for no one. They believe that, as moral agents, we need only be concerned with the effects of our acts on all of the people who are ever actual. We need not consider people who are merely possible—those who never do live but merely might have lived. On this view, the point about identity makes a great moral difference. The effects of our two choices, the predictable catastrophe, and the decline in the quality of life, can be morally totally ignored.

We hoped to resolve this disagreement by appeal to a familiar principle. I suggest now that this cannot be done. To criticize our choice, we must appeal to a claim like (A). And we have yet to explain why (A) should have any weight. To those who reject (A), we do not yet have an adequate reply.

To explain (A), and decide its weight, we would need to go deep into moral theory. And we would need to consider cases where, in the different outcomes of our acts or policies, different numbers of people would exist. This is much too large a task to be attempted here.

I shall therefore end with a practical question. When we are discussing social policies, should we ignore the point about personal identity? Should we allow ourselves to say that a choice like that of the Risky Policy, or of Depletion, might be against the interests of people in the further future? This is not true. Should we pretend that it is? Should we let other people go on thinking that it is?

If you share my intuitions, this seems permissible. We can then use such claims as a convenient form of short-hand. Though the claims are false, we believe that this makes no moral difference. So the claims are not seriously misleading.

Suppose instead that you do not share my intuitions. You believe that, if our choice of Depletion would be worse for no one, this must make a moral difference. It would then be dishonest to conceal the

point about identity. But this is what, with your intuitions, I would be tempted to do. I would not want people to conclude that we can be less concerned about the more remote effects of our social policies. So I would be tempted to suppress the argument for this conclusion.

Theoretical Footnote: How might the attempt to justify claim (A) take us far into moral theory? Here are some brief remarks. Consider any choice between two outcomes. Figure 10.2 shows that there are three kinds of choice. These can be distinguished if we ask two questions: “Would all and only the same people ever live in both outcomes?” “Would the same number of people ever live in both outcomes?”

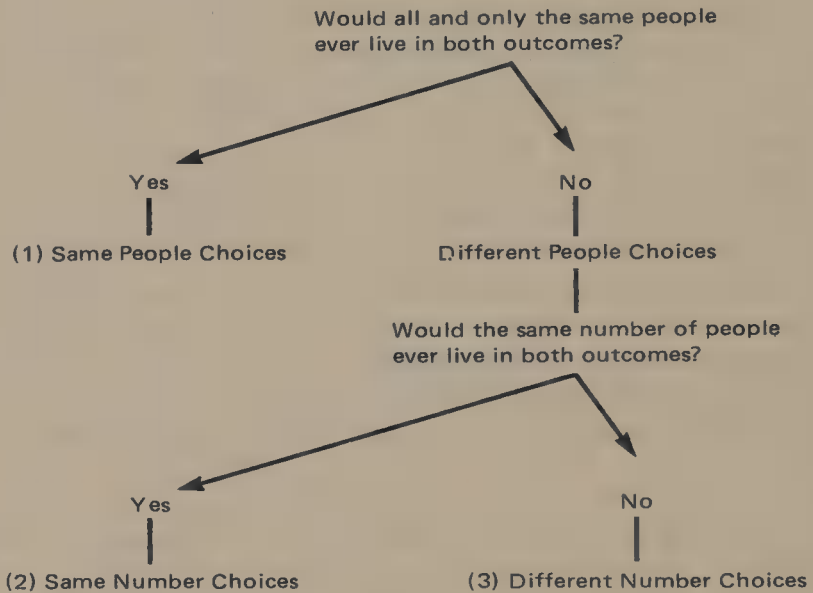


Figure 10.2. Effects of Choice Between Two Outcomes

Of these three types of choice, it is the first and third that are important. Most of our moral thinking concerns Same People Choices, where there is a given group of people whom our acts may affect. We seldom consider Different Number Choices. Those who do have found them puzzling. What this essay has discussed are the second group, Same Number Choices. These are much less puzzling than Different Number Choices. But they are not common. Once we have moved outside Same People Choices—once we are considering acts that would cause different people to exist—it is seldom true that in all of the relevant outcomes the very same numbers would exist.

According to claim (A), it is bad if those who live are worse off than those who might have lived. This claim applies straight-

forwardly only to Same Number Choices. Can we extend (A) to cover Different Number Choices? One extension would be the so-called "Average View." On this view, it would be worse for there to be more people if the average person would be worse off. The Average View, though popular, can be shown to be implausible.³ But this does not cast doubt on (A). What it shows is that (A) should not be thought to cover Different Number Choices. We should restate (A) to make this explicit. But (A) can be made to cover Same People Choices. Our restatement might be this:

(B) If the same number of lives would be lived either way, it would be bad if people are worse off than people might have been.

The two occurrences of "people" here may refer to *different* people. That is how (B) can cover Same Number Choices. But it can also cover Same People Choices. (B) here implies that it is bad if people are worse off than *they* might have been.

Now consider a more familiar principle. This appeals to the interests of those whom our acts affect. One statement might be this:

The *Person-Affecting Principle*, or PAP: It is bad if people are affected for the worse.

What is the relation between (B) and the PAP?⁴ In Same People Choices, these claims coincide. If people are worse off than they might have been, they are affected for the worse. So it will make no difference whether we appeal to (B) or to the PAP.⁵

The two claims diverge only in Same Number Choices. These are what my essay has discussed. Suppose that you share my intuitions, thinking that the point about identity makes no moral difference. You then believe that in Same Number Choices we should appeal to (B) *rather than* the PAP. If we choose Depletion, this will lower the quality of life in the further future. According to (B), this is a bad effect. When we see the point about identity, we see that this effect will be worse for no one. So it is not bad according to the PAP. If we believe that the effect is just as bad, we will here have no use for the PAP. Similar remarks apply to the choice between the two medical programs. If we believe these two programs to be equally worthwhile, we shall again appeal to (B). We shall have no use for the PAP. It draws a moral distinction where, in our view, no distinction should be drawn. It is thus like the claim that it is wrong to enslave whites.

To draw these remarks together: in Same People Choices, (B) and the PAP coincide. In Same Number Choices, we accept (B) rather than the PAP. So, wherever the claims diverge, we prefer (B).

There remain the Different Number Choices. Since we have restricted (B), we shall need some wider claim to cover these. Call this claim (X). I am not sure what (X) should be. But, if you have shared my intuitions, we can expect this. We shall have no further use for

(B). It will be implied by (X).⁶ So we can expect (X) to inherit (B)'s relations to the PAP. Wherever the claims diverge, we will prefer (X). In Same People Choices, (X) will imply the PAP. It will here make no difference to which we appeal. These are the cases with which most moral thinking is concerned. This explains the reputation of the PAP. This part of morality, the part concerned with human welfare, is usually thought of in person-affecting terms. We appeal to the interests of those whom our acts affect. Even after we have found (X), we may continue to use the PAP in most cases. But it will be only a convenient form of short-hand. In some cases, (X) and the PAP will diverge. And we will here appeal to (X) rather than the PAP. We will here believe that, if an effect is bad according to (X), it makes no moral difference whether it is also worse for any particular people. The PAP draws a distinction where, in our view, no distinction should be drawn. We may thus conclude that this part of morality, the part concerned with human welfare, cannot be explained in person-affecting terms. Its fundamental principle will not be concerned with whether acts will be good or bad for those people whom they affect. If this is so, many moral theories need to be revised.⁷

Notes

1. The first third of this section is adapted from my "Future Generations: Further Problems," *Philosophy & Public Affairs* 11, no. 2 (Spring 1982).

2. Thus we might say: "We are benefited only if the alternative would not have been worse for us." If we had never existed, this would not have been worse for us." These and similar arguments I claim not to be decisive in my "Future Generations." Even if it can be in our interests to have been conceived, most of my later claims would still stand.

3. See my "Future Generations," section IX, and Jefferson McMahan's "Problems of Population Theory" in *Ethics* (October 1981).

4. On the assumption that it cannot be in or against our interests to have been conceived. If we drop this assumption, some of the following claims need to be revised. Again, see my "Future Generations."

5. Does the equivalence go the other way? If people are affected for the worse, does this make them worse off? There is at least one exception: when they are killed. (B) should be revised to cover such exceptions. Only this ensures that, in Same People Choices, B and the PAP always coincide.

6. Consider the best-known candidates for the role of (X): the Average and Total Views. In their hedonistic forms, the Average View calls for the greatest net sum of happiness per life lived, the Total View simply calls for the greatest total net sum of happiness. When applied to population policy, these two views lie at opposite extremes. But when applied to Same Number Choices, both imply the hedonistic form of (B). This suggests that, whatever (X) should be, it, too, will imply (B). The difference between the candidates for (X) will be confined to Different Number Choices. This would be like the fact that only in Same Number Choices does (B) diverge from the PAP. I shall discuss these points more fully in my book *Reasons and Persons*, Oxford University Press, 1984.

7. We can expect that we will also change our view about certain common cases (one example might be abortion). But most of our moral thinking would be unchanged. Many significant relations hold only between particular people. These include, for instance, promising, friendship, and (if we are politicians) representation.

My remarks do not apply to these special relations, or to the obligations which they produce. My remarks apply only to our general obligations to benefit and not to harm. Since they apply only to these obligations, and they make a difference only when we can affect who will later live, my conclusion may seem overstated. But consider a (grandiose) analogy. In ordinary cases, we can accept Newton's Laws. But not in all cases. And we now believe a different theory.