

DEREK PARFIT

# Future People, the Non-Identity Problem, and Person-Affecting Principles

## I. FUTURE PEOPLE

Suppose we discover how we could live for a thousand years, but in a way that made us unable to have children. Everyone chooses to live these long lives. After we all die, human history ends, since there would be no future people. Would that be bad? Would we have acted wrongly?

Some pessimists would answer No. These people are saddened by the suffering in most people's lives, and they believe it would be wrong to inflict such suffering on others by having children. In earlier centuries, this bleak view was fairly plausible. But our successors would be able to prevent most human suffering.

Some optimists would also answer No. These people believe that most people's lives are worth living. But they accept two *Strong Narrow Person-Affecting Principles*. On the *Narrow Telic Principle*:

One of two outcomes cannot be worse if this outcome would be worse for no one.

*Editor's Note.* Derek Parfit submitted a draft of this essay to *Philosophy & Public Affairs* on January 1, 2017, in a state that he regarded as unfinished. Parfit died unexpectedly in the early hours of the morning on January 2. The draft contained no notes, and it had various typographical errors and other small mistakes that Parfit would certainly have corrected before publication. The word "[unfinished]" appears at the end of the draft. In a cover note, Parfit indicated that the submitted draft was missing a conclusion and that he hoped to improve the article further with the aid of reviewer comments. After being accepted, the draft was sent to Jeff McMahan, who bears primary responsibility for editing Parfit's manuscripts for posthumous publication and with whom Parfit had discussed successive drafts of the article. McMahan supplied bibliographical citations, made minimal changes to correct typographical and other minor but potentially distracting errors in the text, and inserted a few notes. He has done so in close consultation with Larry Temkin, with whose work Parfit is in dialogue at various points in the article, and made alterations only with Temkin's agreement. The editors are grateful to McMahan and Temkin for their invaluable assistance in bringing this article to its final published state.

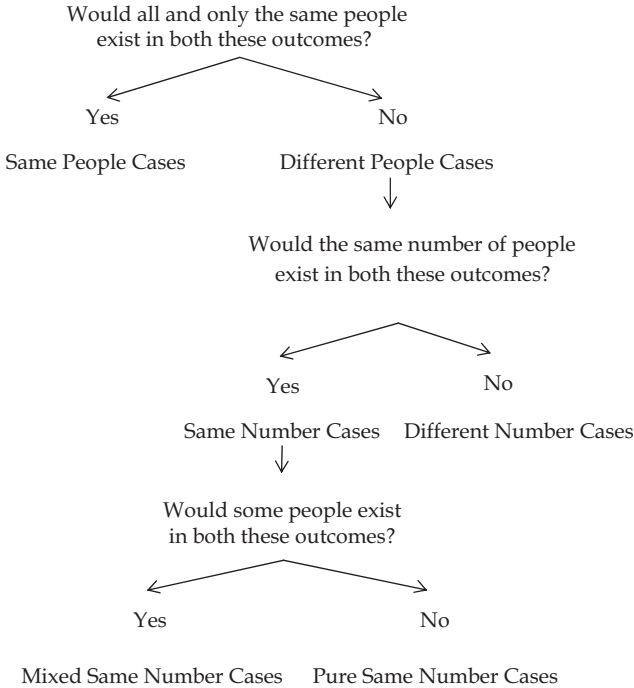
*On the Narrow Deontic Principle:*

An act cannot be wrong if this act would be worse for no one.

It would not be worse, these principles imply, if there were no people, since there would be no one for whom that would be worse. Nor would we be acting wrongly if we all chose to have no children, thereby ending human history.

These principles are, I believe, deeply mistaken. Given what our successors could achieve in the next million or billion years, here and elsewhere in our galaxy, it would be likely to be very much worse if there were no future people. But these principles are not obviously mistaken. We may doubt that anything could be bad if it would be bad for no one.

When we compare two outcomes, or ways in which things might go, there are several possibilities. We can ask



Different Number Cases raise the hardest questions. One such question is

Q1: Would the existence of more people be in itself better, if these people's lives would be worth living?

On the Narrow Telic Principle, the answer is No. It could not be in itself better if more such people existed because it would not have been worse for these people if they had never existed. Some of us find this answer plausible. Others believe the answer to be Yes, and others are undecided, or have no view.

It may help to think first about Pure Same Number Cases. We can ask

Q2: Compared with the existence of some people whose lives would be worth living, would it be in itself better if there existed instead the same number of other people whose lives would be more worth living?

As before, on the Narrow Telic Principle, the answer is No. It could not be better if these other people existed because these people's non-existence would not have been worse for them. When applied to Q2, this answer is less plausible. We may doubt that it would be in itself better if more people existed. But most of us would believe that

(A) *if* more people existed, it would be in itself better if the people who existed were the ones whose lives would be more worth living. This outcome would be better even though these people's non-existence would not have been worse for them.

In believing (A), we would be rejecting the Narrow Telic Principle.

We have other reasons to start by considering Same Number Cases. Many of our acts will indirectly affect the number of people who will later exist. But in most cases, we cannot predict whether our acts would increase or reduce this number. We can justifiably ignore such unpredictable effects. Many of our acts will also indirectly affect the identity of future people, or *who* are the people who will later exist. That our acts will have such effects is much easier to predict. This fact raises problems that, until fairly recently, were overlooked.

To introduce these problems, we can first compare three imagined cases. In our first Same People Case, we can suppose that

*Ruth*, who is pregnant, knows that, unless she takes some painless treatment, the child she is carrying would have some disease that would kill this child at the age of forty. If Ruth takes this treatment, this child would live to eighty.

It would clearly be wrong for Ruth to refuse to take this treatment, since that would be much worse for her child. Suppose next that

*Sarah* must decide whether to have a child. Sarah knows that any child whom she conceives would have this same disease, and would live to only forty. She also knows that, because this disease would have no earlier effects, any such child's life would be likely to be well worth living. Sarah and her husband strongly want to have a child, and there is no existing child whom they could adopt.

As a Different Number Case, this raises harder questions. Most of us would believe that it would not be wrong for Sarah to have such a child. On this view, we can justifiably have children if we would love them and we can justifiably believe that their lives would be likely to be well worth living. We don't need to ask whether it would be in itself better if these children existed.

As our first Same Number Case, we can suppose that

*Clare* knows that, if she conceives some child now, this child would have this same disease, and would live to only forty. If Clare waits for two months, she would later conceive a child who would not have this disease, and would live to eighty.

This case challenges the Strong Narrow Person-Affecting Principles. Most of us would believe that

(B) it would be worse if Clare conceives a child now who would live to only forty. She ought to wait and conceive a child who would live to eighty.

But if Clare conceives a child now, that would not be worse for this child. This child's life would be likely to be well worth living, and if Clare had waited, this child would never have existed. It would have been a *different* child whom Clare would have later conceived and who would have lived to eighty. The Strong Narrow Person-Affecting Principles therefore conflict with (B), since these principles imply that it would not be worse if Clare conceives a child now, nor would this act be wrong. If we believe (B), we must reject the Strong Narrow Person-Affecting Principles and defend our beliefs in some other way. Since this problem arises when in different possible outcomes different people would exist, I called this *the Non-Identity Problem*.

Here is one way in which this problem has been overlooked. We could truly claim that, if Clare conceives her child now, that would be worse for her child. But in the sense in which this claim is true, the phrase "her child" does not refer to one particular person. This phrase refers to any future person who could later be truly called Clare's child. If Clare conceives a child now, who would live to only forty, that wouldn't be worse for the actual person who would be Clare's child. If Clare had waited, this person would never have existed, since it would have been a different person who could have later been truly called Clare's child.

This problem can arise in other ways, and on a different scale. Given the facts about human reproduction, it is true of most people that, if their parents' lives had gone even slightly differently before these people were conceived, these people would not have been conceived, and their parents would have had different children. Most of our choices between two acts or policies would affect the details of our own and other people's lives, in ways that would cause different future people to exist. These effects would spread, and would not, like ripples in a pool, diminish, so that in the further future two quite different sets of people would exist. The Strong Narrow Principles cannot be plausibly applied to such cases.

Suppose, for example, that we and the other members of some large community could choose between two energy policies, one of which would be cheaper but would increase global warming, thereby having various effects that would greatly lower the quality of life that would be had by very many people in several later centuries. Some of the effects of our policy—such as floods, droughts, heat waves, and hurricanes—would kill many of these future people. Despite having these effects,

our choice of this energy policy would not be worse for any of these people, not even those who would be killed. If we had chosen the more expensive policy, which would not have had these bad effects, these future people would never have existed, and their non-existence would not have been better for them. It would have been different people who would have existed instead, and lived much better lives. Since our choice of the cheaper policy would not be worse for any of these future people, the Narrow Principles imply that this choice would not make things go worse, and could not be wrong. These implications seem to me and others to be clearly false, and to give us decisive reasons to reject these Narrow Principles. These choices and acts would make things go much worse, and would be wrong. It would make no moral difference, I believe, that these choices and acts would be worse for no one.

When I first thought about the Non-Identity Problem, I assumed that most people would accept this *No Difference View*.<sup>1</sup> That is not yet true. Of those who have thought about this problem, some still accept the Strong Narrow Person-Affecting Principles. These people continue to believe that our choices or acts could not make things go worse, or be wrong, if these choices or acts would be worse for no one. Many other people accept what we can call a *Two-Tier View*. On such views, we have some moral reasons not to act in ways that would have such bad effects on many future people. But these reasons are weaker than they would have been if these bad effects had been worse for these people.

The Non-Identity Problem must be either practically or theoretically important. If the No Difference View is false, this problem is practically important, since it would matter less whether our acts or policies would have these bad effects. If the No Difference View is true, this problem is theoretically important, since many moral theories imply or assume that this view cannot be true.

When I first defended the No Difference View, I made what I now believe to be a bad mistake. I suggested that, when we consider the cases that raise the Non-Identity Problem, we should appeal to principles that are *impersonal* in the sense that they do not appeal to facts

1. See Derek Parfit, *Reasons and Persons*, reprinted with corrections (Oxford: Oxford University Press, 1985), p. 366.

about what would affect particular people for better or worse.<sup>2</sup> These principles claim that some outcomes would be worse than others even though these outcomes would be worse for no one. My suggested principle was

Q: If in two outcomes the same number of people would later exist, it would be worse if the people who existed would be people whose quality of life would be lower.<sup>3</sup>

This principle applies whether or not, in these outcomes, the same people would exist. Q, I admitted, needs to be derived from some wider theory that can also be plausibly applied to Different Number Cases. Many people reject impersonal principles like Q because they believe that this part of morality—the part concerned with effects on well-being—ought to be explained in person-affecting terms.

I suggested how we could give such an explanation. Though we ought to reject the Narrow Person-Affecting Principles, we might appeal instead to what I called *Wide Person-Affecting Principles*. But I then mistakenly rejected these Wide Principles because they seemed to imply

*The Repugnant Conclusion*: Compared with the existence of many people whose quality of life would be very high, there is some much larger number of people whose existence would be better, even though these people's lives would be barely worth living.

I now believe that, as I shall argue below, one of these Wide Principles both provides the best solution to the Non-Identity Problem and helps us to avoid such repugnant conclusions.

## II. POSSIBLE SOLUTIONS

Many people have proposed other solutions to the Non-Identity Problem. When these people discuss some of the cases that raise this problem, they claim that

2. *Ibid.*, p. 378.

3. *Ibid.*, p. 360.

(C) there are some ways of treating people that would be wrong even if these acts would not, on the whole, be worse for these people. Such acts would wrongly harm people in some particular way, or violate these people's rights, or treat them as a mere means, or fail to respect them, or be condemned by some other moral principle.

Though some of these claims are plausible, they could not, I believe, solve the Non-Identity Problem.

To explain why these proposals fail, we can compare two of my imagined cases. Sarah's case does not raise the Non-Identity Problem. Since any child whom Sarah conceives would live to only forty, Sarah has no reason to choose between the possible children whom she might conceive. Sarah must decide whether to have any child or no child. She also knows that any child of hers would have a life that, despite being fairly short, would be likely to be well worth living. Given these facts, as I have said, most of us would believe that if Sarah knowingly conceives such a child she would not be acting wrongly. Sarah's act would not harm her child, or violate this child's rights. Sarah's act might have been wrong if any child of hers would have had some much greater disadvantage, which would have made this child's life doubtfully worth living. We might then have plausibly appealed to such a child's rights, or to the harm that Sarah's act would have imposed on this child. But since this child's disadvantage is that of living to only forty, Sarah's treatment of this child is not wrong.

Suppose next that, like Sarah, Clare knowingly conceives a child who will live to only forty. Clare also knows that this child's life would be likely to be well worth living. If we believe that Clare's act is wrong, we could not defend this belief by appealing to claims like those mentioned in (C). We could not claim that Clare's act is wrong because this act would be worse for this child, or would harm this child, or violate any of this child's rights. Clare treats her child in the same way in which Sarah treats her child, and Sarah's act is not wrong. If Clare's act is wrong, as most of us would believe, this act cannot be made to be wrong by this act's effects on Clare's child, or on other existing people. This act is wrong because, unlike Sarah, Clare could have easily conceived a different child who would have lived to eighty.

When people respond to the Non-Identity Problem by making claims like those mentioned in (C), most of these people seem to



assume that an act's wrongness depends only on how this act would affect some of the people who are or will be actual. In the cases that raise the Non-Identity Problem, that is not true. An act's wrongness may depend on how some other possible act would have affected people who are merely possible, but who, if we had acted in this other way, would have been actual. To solve the Non-Identity Problem, we must explain how certain acts can be made to be wrong by such facts about some merely possible people, even if these acts would not be worse for any actual people.

In an excellent recent book, David Boonin forcefully criticizes several of these proposed solutions to the Non-Identity Problem.<sup>4</sup> Boonin then proposes his own solution. Boonin rejects the Strong Narrow *Telic* Principle since he believes that, in the cases that we are considering, certain acts or policies would make things go much worse, even though these acts or policies would be worse for no one. But Boonin accepts the Strong Narrow *Deontic* Principle. He claims that

(D) there is no Non-Identity Problem. In all of the cases that seem to raise this problem, we can justifiably appeal to this Narrow Deontic Principle. We can claim that, if certain acts would be worse for no one, these acts cannot be wrong.<sup>5</sup>

Boonin admits that, in some cases, this Deontic Principle has implausible implications. Boonin imagines someone, *Wilma*, who acts like my imagined Clare. To avoid the inconvenience of taking some pills, Wilma chooses to conceive a child who will have some genetic disease or disability rather than conceiving some other, healthy child. It would be implausible, Boonin claims, to deny that Wilma's act is wrong. Boonin also supposes that,

as members of some wealthy society, we could choose some *Risky Energy Policy* that would make things go slightly better for us than some safe alternative, but would also be likely to cause toxic

4. David Boonin, *The Non-Identity Problem and the Ethics of Future People* (New York: Oxford University Press, 2014).

5. This is not a quotation from Boonin's book but is Parfit's understanding of the main claim underlying Boonin's position.

wastes to leak five hundred years later, killing tens of thousands of people.<sup>6</sup>

Our choice of the Risky Policy would not be worse for the people who would later be killed, since if we had chosen some safe alternative these people would never have existed, and that would not have been better for them. But we may find it hard to believe that, as the Narrow Deontic Principle implies, our choice of this risky policy would not be wrong. As Boonin writes:

Will we really be willing to accept that it is not wrong to kill tens of thousands of innocent people just in order to enjoy a slightly higher standard of living?<sup>7</sup>

Boonin calls this the *Even More Implausible Conclusion*.

Since these conclusions are implausible, we cannot solve the Non-Identity Problem merely by accepting them. Boonin's proposed solution makes some further claims. These conclusions, Boonin argues, are not as implausible as they seem to be. There are other cases in which, though we could make things go much better by acting in some way at little cost to ourselves, failing to act in this way would not be wrong. In Boonin's main example, we rich people know that if we gave some fairly small sum of money to some aid agency, we would enable this agency to save some distant child's life.<sup>8</sup> Most of us would believe that, though it would be better if we gave this money to this aid agency, our failure to give this money would not be wrong. Similar remarks apply, Boonin claims, to Wilma's act of conceiving a diseased child rather than a different, healthy child. Boonin calls it "perfectly plausible" that, like our failure to give the money that would save some distant child's life, Wilma's act would not be wrong. We are not always morally required to do what would prevent things from going much worse. The case of the

6. This indented passage is not a quotation from Boonin but is a new statement of an example first presented in *Reasons and Persons* (pp. 371–77), a slightly modified version of which is discussed in various places in Boonin's *Non-Identity Problem*. The version here combines details from both the original version and Boonin's.

7. Boonin, *The Non-Identity Problem*, p. 220.

8. Boonin does not regard this as his "main example" and wrote to Parfit in explanation; but this was so soon before Parfit died that he was unable to make further revisions.

Risky Energy Policy, Boonin concedes, is more troubling. It is hard to believe that we could justifiably choose some policy that would later kill tens of thousands of people just in order to enjoy a slightly higher standard of living. But Boonin claims that there is no other, less implausible solution to the Non-Identity Problem.

That is not, I believe, true. Since Boonin accepts the Strong Narrow Deontic Principle, his argument for his solution assumes that

(E) in the cases that seem to raise the Non-Identity Problem, if some act would be worse for no one, this act cannot be wrong.

We can justifiably reject (E). We can claim that

(F) it is wrong to do what would make things go much worse if some other possible act would not be worse for anyone, or violate anyone's rights, or have any other moral flaw.

Boonin writes that, if Wilma chooses to conceive a diseased child now rather than later conceiving a different, healthy child, Wilma's act would be "a little bit better for her but far worse from the point of view of overall human well-being."<sup>9</sup> We can plausibly assume that, if Wilma instead conceived a healthy child, this other act would not be worse for Wilma or for anyone else, nor would this act violate anyone's rights, or have any other moral flaw. As (F) would then imply, Wilma's act of conceiving a diseased child would be wrong. Similar claims apply to the other cases that raise the Non-Identity Problem. In such cases, (E) and (F) conflict, and (F) is more plausible.

Boonin might reply that, in such cases, there would seldom be some other possible act of the kind described by (F). In most cases of these kinds, he might say, we could avoid making things go much worse only at some significant cost to ourselves. We can suppose that, in a different version of Wilma's case, Wilma could conceive a healthy child only by acting in some way that would impose some significant burden on herself. This would be like a case in which it is only by giving much money that we could enable some aid agency to save some distant child's life. Boonin might then claim that, just

9. Boonin, *The Non-Identity Problem*, p. 195.

as it would not be wrong for us to choose not to give this much money to this aid agency, it would not be wrong for Wilma to choose not to bear the burden that would enable her to conceive a healthy child.

This reply could not, however, defend Boonin's (E). We are not here discussing the demandingness of morality, and asking how great the burdens are that we ought to bear when that is our only way to avoid making things go much worse. We are asking whether, in the cases that raise the Non-Identity Problem, some act could be wrong even though this act would not be worse for anyone. This question takes its clearest form when we suppose that such an act would make things go much worse and that some other possible act would not be worse for anyone, or violate anyone's rights, or have any other moral flaw. Given these facts, as (F) implies, the first of these acts would be wrong. These cases are enough to show that (E) is false. Boonin's claims do not, I conclude, solve the Non-Identity Problem.

There is a better solution. Following McMahan, we can claim:

If someone is caused to exist and to have a life that is worth living, that is good for this person, giving him or her an *existential* benefit. There are similar existential harms.<sup>10</sup>

We can next distinguish two other person-affecting principles. According to

*The Weak Narrow Principle:* One of two outcomes would be in one way worse if this outcome would be worse for people.

According to

*The Wide Principle:* One of two outcomes would be in one way worse if this outcome would be less good for people, by benefiting people less than the other outcome would have benefited people.

10. Jeff McMahan, "Causing People to Exist and Saving People's Lives," *Journal of Ethics* 17 (2013): 6–7.

When we are considering outcomes in which all of the same people would exist, these Narrow and Wide Principles coincide. In such cases, if one of two outcomes would be *worse* for some people, this outcome would also be *less good* for these people. But when some people would exist in only one of two outcomes, these principles may not coincide. Even if some outcome would be worse for no one, this outcome may be less good for people than the other outcome would have been for the different people who, in this other outcome, would have existed. That would be true when this other outcome would have given these other people greater existential benefits. In such cases, these Narrow and Wide Principles may deeply conflict.

We ought, I shall argue, to reject the Weak Narrow Principle. We can claim that there are existential benefits, and appeal to the Wide Principle, thereby solving the Non-Identity Problem.

Suppose again that, by choosing the cheaper energy policy that would increase global warming, we would greatly lower the quality of life of very many future people, and would indirectly cause many of these people to be killed. Our choice of this policy would not be worse for these future people, since it would not have been better for them if they had never existed. In causing these people to exist, our policy may even be good for them. But this policy would be much less good for these people than the more expensive, better policy would have been for the different people who would have existed instead and had much better lives. That is how our policy would make things go worse, and would be wrong. It is similarly true that, when Clare conceives a child who will live to only forty, Clare's act is not worse for this child, but Clare's act benefits this child much less than Clare could have benefited a different child who would have lived to eighty. That is how Clare's act would make things go worse, and would be wrong.

Return next to my imagined case in which we discover how we could live for a thousand years, but in a way that would make us unable to have children. If we all chose to benefit ourselves in this way, we would end human history. The Strong Narrow Principles imply that, if these acts would be worse for no one, they could not make things go worse, and would not be wrong. We should reject these claims. These acts would be likely to make things go very much worse by being very much less good for us than the survival of humanity would be for very many

future people. It would be similarly wrong to increase or ignore existential risks to the future of humanity.

### III. EXISTENTIAL BENEFITS AND HARMS

In defending these claims, we can first ask

Q3: If our lives are worth living, were we benefited by being caused to exist?

Many people would answer No. Some of these people argue:

Benefits are comparative. One of two outcomes benefits us only when the other outcome would have been worse for us.

If we had not been caused to exist, this fact could not have been worse for us, since there would have been no us for whom our non-existence could have been worse.

Therefore,

Being caused to exist cannot have benefited us.

We can reject this argument. We can reply:

Benefits are intrinsic. Our being caused to exist benefited us by being intrinsically good for us, even though our non-existence would not have been worse for us.

We can add:

It is true that, if we had never existed, there would have been no us who did not receive this benefit. But there *is* an actual person, *us*, who did.

Things of some kind are *intrinsically* or *in themselves* good or bad when what makes them good or bad are their intrinsic properties, or what they are like in themselves. We can consider each of these things

on their own, and ask how good or bad this thing is. There are some other things whose goodness or badness is not intrinsic, since it depends on how these things are related to other things. One trivial example is the goodness of sporting teams. We cannot consider each team separately and ask how good it is. To assess the goodness of these teams, we must know what happens when these teams play other teams. Some team would be good, for example, if this team beats most other teams. These kinds of goodness or badness we can call *essentially comparative*.<sup>11</sup>

When people give the argument stated above, they rightly assume that most benefits are in one way comparative, since most of the outcomes that benefit us are better for us than some alternative. Having friends, for example, is better than having no friends. But the goodness of having friends is not *essentially* comparative. Friendships are intrinsically good, as are most of the other good things in our lives. We can therefore reject the first premise of the argument stated above. We can be benefited by being caused to exist and to have a good and happy life, even though the alternative, in which we never existed, would not have been worse for us. These are the benefits that I am calling existential.

Similar claims apply more clearly to existential harms. Some lives are intrinsically bad, and worse than lives that are merely not worth living. Suppose that, in *Case One*, either

Sam will exist and have fifty years of suffering

or

Sam will never exist.

It would be worse if Sam was caused to exist and to live this wretched life. If we believed that all benefits and harms are essentially comparative, and we appealed only to the Weak Narrow Principle, we could not explain why it would be worse if Sam was caused to have this wretched

11. See Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (New York: Oxford University Press, 2012), pp. 229–31.

life. If Sam had never existed, there would not have been a Sam for whom his non-existence would have been better.

Some people suggest that, to explain why it would be worse if Sam was caused to have such a life, we would have to appeal to some impersonal principle, such as the claim that it would be worse if there was more suffering. But we don't need to appeal to any such principle. We can claim that most harms are intrinsic. The badness of these harms does not consist in their being in some way worse than some alternatives.

Here is another way to explain this distinction. If we assumed that harms and benefits are all essentially comparative, we might claim that

(G) it would be worse if Sam exists and has this wretched life because it would have been better for Sam if he had never existed.

But this claim misleadingly suggests that, if Sam had never existed, there would have been a person, Sam, for whom his never existing would have been better. That seems close to a contradiction. We might distinguish between *being* and *existing*, so that we could claim that Sam would have *been* a non-existing person.

We can, I believe, draw such distinctions. But there is no need to do that here. We can claim instead that

(H) it would be worse if Sam exists and has this wretched life because Sam's life would be *bad* for him, and his non-existence would not have been *bad* for him.

In claiming (H), we are not comparing Sam's state, if he exists, with the state Sam would have been in if he had never existed. We are merely claiming that being in a state that is bad is worse than *not* being in any bad state. Having fifty years of suffering is intrinsically bad, and such badness doesn't have to involve being worse than some other state that either Sam, or we, might have been in. All the badness of Sam's suffering would be had by this suffering. That is what makes such badness intrinsic rather than comparative.

As we have seen, there are two closely related questions here. One is the question whether our states can be good or bad only by being better or worse than other states. The other is the question whether it can



be good, or bad, to be caused to exist and to be in certain states, rather than never to exist. To illustrate this distinction, we can appeal to similar claims about the badness or goodness of being killed. We can agree that, in most cases,

(I) we are harmed only if we are caused to be worse off than we would otherwise have been.

But this claim does not apply to our being killed, since having ceased to exist is not a way of being badly off. We are harmed by being killed, we could instead claim, when and because the rest of our life would have been intrinsically good for us.

Similar claims apply to our being caused to exist. We can agree that in most cases,

(J) we are harmed only if some relevant alternative would have been better for us.

But this claim does not apply to our being caused to exist. Like having ceased to exist, never existing is not a way of being badly off. Unlike our being killed, our never existing could not even be bad for us, or good for us, since there would have been no *us* for whom our non-existence could have been good, or bad. But these facts do not imply that it could not be bad for us to be caused to exist and have a life that is full of suffering. If we exist, there *is* an *us* for whom our having being caused to have this life was bad. Nor would our life be made to be less bad for us by the fact that our non-existence would not have been better for us. Similar claims apply to lives that are intrinsically good. We can be benefited by being caused to exist and to have an intrinsically good life, though our non-existence would not have been worse for us.

Here is another way to defend these claims. It is clear that

(K) if Sam already exists, and we acted in some way that caused Sam's life to be full of suffering, we would be harming Sam by doing something that would be bad for him.

If we accept (K), we should also accept that

(L) if we acted in some way that caused Sam's life to be full of suffering, we would be harming Sam by doing something that would be bad for him.

These acts would both be ways of causing Sam to have a life that was full of suffering. It would not be less bad for Sam, nor would it do less to make the outcome worse, if we caused Sam to have this wretched life by causing Sam to exist.

Some defenders of the Weak Narrow Principle accept some of these claims. These people believe that it can be bad for someone to be caused to exist and have a life that is intrinsically bad, even though the alternative would not have been better for this person. But though these people accept that there can be such existential harms, they deny that there can be existential *benefits*. On this asymmetrical view, though it can be *bad* for people to be caused to exist, it cannot be *good* for people to be caused to exist.

This view could not, I believe, be true. It is fairly plausible to claim that all harms and benefits are comparative, since no outcome could harm or benefit us unless some alternative would have been better for us, or worse for us. But we cannot plausibly claim that, though harms are intrinsic, benefits are essentially comparative. If it can be bad for people to be caused to exist and have some wretched life, even though their non-existence would not have been better for them, it could be similarly good for people to be caused to exist and have some happy life even though their non-existence would not have been worse for them. We can at most claim that our moral reasons not to harm people are in some ways stronger than our reasons to benefit people. These differences would be a matter of degree.

When people deny that there are existential benefits, some of them appeal to Jan Narveson's distinction between two ways of increasing happiness. Narveson claimed that, though it is good to make people happy, it is not good but morally neutral to make happy people.<sup>12</sup> We could not defensibly make such a claim about two ways of increasing suffering. Compared with making people miserable, it would be just as bad to make miserable people.

12. Jan Narveson, "Moral Problems of Population," *Monist* 57 (1973): 80.

Suppose next that we reject this asymmetrical view, since we believe that there can be existential benefits as well as existential harms. Some people appeal to another asymmetry. On this view, if we knowingly caused someone to exist whose life would be full of suffering, and worse than nothing, there would be an actual person whom we had harmed, and thereby wronged. If instead we fail to cause some happy person to exist, there would be no actual person whom we had failed to benefit, and thereby wronged. On what we can call this

*No Complainants Claim:* An act cannot be wrong unless there is or will be someone whom this act has wronged.

This claim is wider than the Strong Narrow Deontic Principle, since this claim allows that people might be wronged by some act even though this act is not worse for them. Some of the clearest cases involve wronging people after they are dead.

This *No Complainants Claim* has considerable psychological force. It also has moral force. On Scanlon's plausible version of contractualism, we act wrongly if our act is disallowed by some principle that no one could reasonably reject.<sup>13</sup> If we accept this view, we might add that no one could reasonably reject some principle unless this principle's acceptance would be in some way worse for them. On this narrow version of Scanlon's view, if we fail to benefit people by failing to cause them to exist, we would have done nothing that would be worse for any particular people, so we could not have acted wrongly. We would not have failed to treat some people in some way that we owed to them, since we can't owe anything to people who never exist.

As Scanlon notes, however, these claims do not cover the whole of morality. Nor do they describe the only way in which we have reasons to want to be related to others. Our aim should not be only to avoid doing what would be worse for other people. It is true that, if we don't cause people to exist, there would be no actual people who would have reasons to regret this failure to benefit them. It is also true, however, that if we cause people to exist and have lives that are worth living, there would be actual people who would have reasons to be grateful

13. T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998).

about this way in which we had benefited them. We may believe that, compared with giving people reasons to be grateful about what we have done to them, it is morally more important not to give people reasons to regret what we have done to them. But this would again be only a difference of degree. Here is another similar claim. To avoid giving people reasons to complain about what we had done to them, it is enough to do nothing that would be bad for these people. We could achieve this moral aim in a purely negative way, by doing nothing. Though non-maleficence is negative, beneficence is positive. We can give people reasons to be grateful only by doing things that are good for these people. And though it is morally important not to act badly, it is also morally important to act well.

I shall return to the No Complainants Claim.<sup>14</sup> But I am now discussing, not which acts would be morally permissible or wrong, but which outcomes would be better or worse.

#### IV. NARROW AND WIDE PERSON-AFFECTING PRINCIPLES

We can next draw some more distinctions. Some relation is *transitive* when it is true that, if this relation holds between X and Y, and between Y and Z, this relation must also hold between X and Z. *Being taller than* is one example. If Tom is taller than Dick, who is taller than Harry, Tom must be taller than Harry. Many relations are not transitive. *Loving someone* is one example. It may be true that Tom loves Dick, who loves Harry, but Tom does not love Harry.

There are, I have claimed, two ways of being good or bad, and better or worse than other things. Things of some kind are *intrinsically* good or bad when their goodness or badness depends only on their intrinsic properties, or what they are like in themselves. When we compare the goodness or badness of such things, we can consider each of them separately and try to decide how good or bad they are. *Intrinsically better than* is a transitive relation. If X is in itself better than Y, which is in itself better than Z, X must be in itself better than Z. There are some other things whose goodness or badness is not intrinsic but *essentially comparative*, because how good or bad they are depends on how they

14. Parfit does not return to the No Complainants Claim, though he does return to the issue of complaints later in section V when he discusses *Case Two*.

are related to other things. Of these ways of being *comparatively better than*, some are transitive, but others are not.

The goodness of weight lifters is intrinsic, since it depends entirely on what are the heaviest weights that, when acting on their own, people are able to lift. When applied to weight lifters, the relation *better than* is transitive, because the relation *heavier than* is transitive. If Tom can lift heavier weights than Dick, who can lift heavier weights than Harry, Tom can lift heavier weights than Harry. The goodness of sporting teams, in contrast, is essentially comparative. In trying to decide how good such teams are, we must know what happens when these teams play other teams. We can then describe some ways in which some teams are better than others. According to

*The Numerical Criterion:* One of two sporting teams is better if this team beats more other teams.

On this criterion, the relation *better than* is transitive. If team X beats more other teams than Y, which beats more other teams than Z, X must beat more other teams than Z. There are some other criteria that are essentially *pairwise* comparative, since they imply that the relative goodness of two things depends entirely on some relation between these two things. According to

*The Always Beats Criterion:* One of two sporting teams is better if, whenever these teams play, the first team wins.

This way of being better than is not transitive. It might be true that team X always beats team Y, which always beats team Z, which always beats team X. Such claims can be true when different teams have different weaknesses and strengths. This criterion would then imply that X is better than Y, which is better than Z, which is better than X. Each of these three teams would be worse than one of the others.

When Temkin discusses these criteria, he suggests that we ought to reject the Always Beats Criterion and accept some version of the Numerical Criterion. I suggest that we regard these criteria as describing two ways in which some sporting teams may be better than others. The Numerical Criterion describes a more important way of being a

better team. But if one team always beats another, there is a clear non-transitive sense in which this team is better. We should not claim that X is in *no* sense better than Y even though, whenever these teams play, it is X that wins. X is better than Y, we could say, in the *always beats* sense.

In his magnificent book *Rethinking the Good*, Temkin sometimes claims that, when applied to outcomes, the relation *all things considered better than* is not transitive.<sup>15</sup> But this claim oversimplifies Temkin's view. Temkin rightly assumes that the phrase "better than" can be used in different senses, which refer to different relations. Temkin argues that, though some outcomes are intrinsically better than others in a sense that is transitive, such claims may apply to only a "severely limited part of the normative realm." I believe that many outcomes, and many other things, are intrinsically good or bad, in ways that make them better or worse than others. Such goodness or badness is not essentially comparative, since it does not consist in being related in certain ways to other things.

Return now to

the Weak Narrow Principle: One of two outcomes would be in one way worse if this outcome would be worse for people,

and

the Wide Principle: One of two outcomes would be in one way worse if this outcome would be less good for people, by benefiting people less than the other outcome would have benefited people.

Given the similarity between the meanings of "worse for" and "less good for," there may seem to be little difference between these principles. But that is not so.

There are some deep differences between these principles. One difference, as I have claimed, is in the range of cases to which these principles apply. When we consider the further future, there are many cases in which no one would exist in more than one of the possible outcomes. In these very many cases, the Strong Narrow Telic Principle

15. Temkin, *Rethinking the Good*, esp. pp. 217–28.

implies that none of these outcomes could be worse than any of the others, and the Weak Narrow Principle fails to imply that any of these outcomes would be worse than any others. The Wide Principle, in contrast, applies to all possible cases.

The Wide Principle makes claims about the intrinsic goodness of different outcomes. On this principle, one of two outcomes would be intrinsically worse if this outcome would *benefit people less*, and these relations are both transitive. If A would benefit people less than B, which would benefit people less than C, outcome A must benefit people less than C. Similar claims apply to *harming people more*.

When we apply the Wide Principle, we can consider each outcome on its own, and ask *how good* or *bad* this outcome would be for people. An outcome's intrinsic goodness does not depend on its relation to other outcomes. When we apply the Weak Narrow Principle, which is not about intrinsic goodness or badness, we cannot consider each outcome on its own. We cannot ask whether some outcome would be *better for people*, or *worse for people*, since these relations are essentially comparative, holding only between different outcomes. There is another relevant distinction. When we are considering outcomes in which all and only the same people would exist, the relation *worse for people* may be transitive. But when some people would exist in only some of the outcomes that we are considering, the relation *worse for people* is not transitive. Even if A would be worse for people than B, which would be worse for people than C, A may not be worse for people than C. A may even be better for people than C.

In comparing the Weak Narrow and Wide Principles, we can apply them to some imagined cases in which certain people might exist and have shorter or longer lives. We can suppose that each extra year of life would be an equal benefit, and that there would be no other relevant differences between these imagined people. We can also ignore various indirect effects. My main claims could be applied to more complicated and realistic cases. In some of my imagined cases, the Weak Narrow Principle coincides with the widely accepted

*Pareto Principle:* One of two outcomes would be worse if this outcome would be worse for some people and better for no one.

Suppose that, in *Case Two*, some possible outcomes are:

A: Tom will live to 60	Dick will live to 80	_____
B: Tom will live to 80	_____	Harry will live to 60
C: _____	Dick will live to 60	Harry will live to 80

I use “\_\_\_\_\_” to mean “will never exist.”

The Narrow Principle here implies that outcome A would be in one way worse than outcome B, since A would be worse than B for Tom and B would be worse than A for no one. B would be similarly worse than C by being worse than C for Harry and better than C for no one. C would be similarly worse than A by being worse than A for Dick and better than A for no one.

We can next suppose that there are no other morally relevant differences between these three outcomes. The Narrow Principle then implies that each of these outcomes would be worse than one of the others. On this view, when applied to such outcomes, *all things considered worse than* would not be a transitive relation. If we were using *worse than* in some sense that is not intrinsic but essentially pairwise comparative, these claims would be coherent, since they would be like the claim that, in the *always beaten by* sense, each of three sporting teams is worse than one of the others. But if we used *worse than* and *better than* in these non-transitive senses, it would be harder for us to reach true beliefs about what we had more reason to prefer, and to do. To give one example, our beliefs would imply that, if we changed the world in a series of ways each of which was a change for the better, these changes might together make the world worse than it was at the start.<sup>16</sup> If that were true, it would be very unclear which of these changes, if any, we had more reason to make.

There is a better view. Rather than asking which of these outcomes would be *worse for people* than others, we should claim that these are intrinsic benefits, and ask instead how *good for people* these outcomes would be. In *Case Two*, each of these outcomes would give one person the intrinsic benefit of living to eighty and give one other person the benefit of living to sixty. As the Wide Principle implies, these three outcomes would be equally good because they would be equally good for people.

16. See Temkin, *Rethinking the Good*, pp. 79–85.



We can next briefly return to the Two-Tier View. On this view, we admit that there are existential benefits, but we claim that these benefits have less weight than ordinary comparative benefits.<sup>17</sup> In *Case Two*, some possible outcomes are:

A: Tom will live to 60	Dick will live to 80	_____
B: Tom will live to 80	_____	Harry will live to 60
C: _____	Dick will live to 60	Harry will live to 80

On the Two-Tier View, the comparative benefit of twenty years that B gives Tom morally outweighs the net existential benefit of twenty years that A gives Dick. (This net benefit is the benefit to Dick minus the benefit to Harry.) The Two-Tier View therefore implies that worse than is not transitive, since this view implies that A would be worse than B, which would be worse than C, which would be worse than A. To avoid these implications, we must appeal instead to the Wide Principle.

Suppose next that, in *Case Three*, some possible outcomes are:

D: Mary will live to 70	Kate will live to 50	_____	_____
E: _____	Kate will live to 60	Ruth will live to 25	_____
F _____	_____	Ruth will live to 30	Jill will live to 10

The Weak Narrow Principle here implies that D would be worse than E, since D would be worse than E for Kate and E would not be worse than D for anyone. E would be similarly worse than F since E would be worse than F for Ruth and F would not be worse than E for anyone. These claims are the opposite of the truth. D would be better than E, which would be better than F. If Mary lives to seventy and Kate lives to fifty, that would be better than if Kate lives to sixty and Ruth lives to twenty-five, which would be better than if Ruth lives to thirty and Jill lives to ten.

Suppose next that these claims use *worse than* in a transitive sense. The Weak Narrow Principle would then imply that D would be worse

17. See McMahan, "Causing People to Exist and Saving People's Lives," pp. 15–20.

than F. That is even more clearly false. We have no reason of any kind to believe that D would be worse than F. If two people live for a total of 120 years, that would not be in any way worse than if these people never exist and two *other* people exist instead and live for a total of only forty years. The first of these outcomes would clearly be much better.

Defenders of the Weak Narrow Principle might reply that they are using *worse than* in some non-transitive sense. They might then claim that, though their view implies that D would be worse than E, which would be worse than F, their view does not imply that D would be worse than F. But this reply would not answer these objections. *Case Three* does nothing to support the view that, in comparing these outcomes, we should use *worse than* in some non-transitive sense. And even if defenders of the Weak Narrow Principle deny that, on their view, D would be worse than F, they must admit that, on their view, D would be worse than E, and E would be worse than F. These claims are both clearly false.

*Case Three* also illustrates how the Weak Narrow Principle goes astray. When we apply this principle to any pair of outcomes, this principle takes into account only what would happen to the people who would exist in *both* these outcomes. On this principle, D would be worse than E because D would give Kate ten fewer years of life. This principle ignores the fact that D would give Mary seventy years of life and that E would give Ruth only twenty-five years. When this principle similarly implies that E would be worse than F, because E would give Ruth five fewer years of life, this principle ignores the fact that E would give Kate sixty years of life and F would give Jill only ten years. It is an obvious mistake to ignore these facts.

The Weak Narrow Principle ignores these facts because this principle assumes that benefits are essentially comparative. On this principle, D is worse than E because D is worse than E for Kate, and D is not better than E for Mary. It is irrelevant that, in outcome D, Mary lives for seventy years. These seventy years of life do not benefit Mary, this principle assumes, because we are benefited only when some alternative would have been worse for us, and it would not be worse for Mary if the actual outcome would be E, in which Mary never exists. We should reject these claims. Mary's seventy years of life in E would be intrinsically good for her. To explain the different claim that F would be worse than D, we can appeal to intrinsic benefits and to the Wide Principle. We can claim that, though F would not be *worse* for people than D, F

would be *less good* for people than D. If Ruth lives to thirty and Jill lives to ten, that would be less good for them, or would benefit them less, than the benefits to Mary and Kate of living to seventy and to fifty.

Suppose next that, in *Case Four*, some possible outcomes are:

A: Tom will have 50 years of pain	_____	Harry will have 1 day of pain
B: _____	Dick will have 1 day of pain	Harry will have 2 days of pain

This case illustrates another objection to the Weak Narrow Principle. Because this principle is essentially comparative, asking only which outcomes would be worse for people, this principle ignores the intrinsic badness of these outcomes. Since B would be worse than A for Harry, and B would not be better than A for anyone, the Weak Narrow Principle implies that B would be worse than A. This implication is absurd. Harry's one extra day of pain in B does very much less to make the outcome worse than Tom's more than eighteen thousand days of pain in A. As the Wide Principle implies, A would be much worse than B.

These three cases provide, I believe, decisive objections to the Weak Narrow Principle. This principle has implications that are clearly false, and it has these implications because it ignores some of the morally relevant facts.

#### V. THE GOODNESS OF OUTCOMES AND THE RIGHTNESS OF ACTS

Temkin questions some of these claims. Return to *Case Two*, in which some possible outcomes are:

A: Tom will live to 60	Dick will live to 80	_____
B: Tom will live to 80	_____	Harry will live to 60
C: _____	Dick will live to 60	Harry will live to 80

The Weak Narrow Principle implies, I have claimed, that A would be worse than B, which would be worse than C, which would be worse than A. That is not true, since these outcomes would be equally good.

Temkin suggests an answer to this objection. If we accept the Weak Narrow Principle, Temkin suggests, we might claim that

(M) if only outcomes A and B are possible, A would be worse than B, because A would be worse than B for one person and better for no one. If instead all three outcomes are possible, these outcomes would be equally good, because they would all be worse for one person than one of the other outcomes, and worse by the same amount.<sup>18</sup>

If this claim were true, *Case Two* would not provide an objection to the Narrow Principle and the Two-Tier View.

The goodness of outcomes could not, I believe, vary in the way described by (M). The word “outcome” is sometimes used to refer only to the effects of some act. I use “outcome” in a wider sense, which refers to any way in which things might go, or might have gone, or to events in a sense that includes acts. If some people suffer, or die young, for example, these would be outcomes in this wide sense. When Temkin supposes that, in *Case Two*, only outcomes A and B are possible, he is supposing that these are the only outcomes that we could bring about. But the goodness or badness of outcomes, or events, does not depend on whether we could bring them about. If ten people’s lives were saved, for example, that would be better than if five people’s lives were saved. This outcome would be better whether or not anyone could save any of these people’s lives. We can similarly claim that, if outcome A would be worse than outcome B, as the Narrow Principle implies, this would be true whether or not we could bring about either or both these outcomes, or could bring about outcome C. Since the goodness of outcomes could not vary in the way described by (M), this defense of the Narrow Principle and the Two-Tier View does not succeed.

Temkin might revise his suggestion. Though the goodness of *outcomes* does not depend on which outcomes are causally possible in the sense that we could bring them about, the rightness of *acts* may depend on which acts are in this sense possible. Suppose that we could either

18. Compare Temkin, *Rethinking the Good*, pp. 427–34.

X: Save the life of some stranger at some great risk to ourselves

or

Y: Do nothing.

If these are the only possible acts, X would be morally admirable, but Y would not be wrong. We are not morally required to risk our own life even if we could thereby save some stranger. Suppose next that we could also

Z: Save both this stranger and another stranger at this same risk to ourselves.

If we knew that Z was possible, our doing X rather than Y would be wrong. If we risk our life to save one person's life, it would be wrong not to save two people's lives at no greater risk to ourselves. Whether doing X would be wrong here depends on which other acts are possible.

Return now to *Case Two*. Temkin suggests that

(N) if we could bring about only outcomes A or B, A would be worse than B, but if we could also bring about outcome C, these three outcomes would be equally good.

The goodness of outcomes could not, I have claimed, vary in this way. But Temkin might instead claim that

(O) if we could bring about only outcomes A or B, we ought to bring about B, but if we could also bring about C, it would not be true that we ought to bring about B, since these three *acts* would then be equally good.

Temkin might explain this claim in a narrow person-affecting way, by appealing to claims about what would be worse for people. Temkin might say that, if we could not bring about C, it would be wrong for us to bring about A, since Tom would have a justified complaint. Tom could appeal to the fact that A would be worse than B for him, and A would not be better than B for anyone else. Since we could not bring about C, neither Harry nor Dick would have a similar complaint. If

instead we *could* bring about C, none of these three people would have a justified complaint against any of these three acts. Though A would be worse than B for Tom and better for no one, B would be similarly worse than C for Harry and better for no one, and C would be similarly worse than A for Dick and better for no one. That is why, if we could bring about each of these outcomes, none of these acts would be what we ought to do. Either these acts would all be permissible, or we ought to choose between them in some other way. Temkin's main claims are about the goodness of outcomes. But Temkin sometimes states these claims in ways that also apply to what we ought to choose and to do. Temkin might therefore be happy to reject (M) and (N) and to appeal instead to (O). If we can reject (M), that would help to defend the view that, when applied to outcomes, *all things considered better than* is a transitive relation. This would be, as Temkin claims, a welcome conclusion.

In distinguishing between the goodness of such outcomes and the rightness of such acts, it may help to consider a simpler case. Suppose that, in *Case Five*, two possible outcomes are:

A: Jack will live to 70	Jane will live to 80	_____
B: Jack will live to 71	_____	Sally will live to 60

Some of our beliefs about such cases may depend on our temporal point of view. Suppose first that none of these three people now exists. Jack, we assume, is a future person in the sense that Jack will exist. Jane and Sally are possible people, either of whom might be a future person. When we ask which of these outcomes would be better, we are asking whether, if one future person would later live for one year longer, that would be better than if the possible person who would later exist would be the person whose life would be twenty years longer. The answer, I believe, is No. If Jack will later live to seventy and Jane will live to eighty, that would be better than if Jack will later live to seventy-one and Sally will live to sixty. On this view, Jack's one extra year of life would do less to make the outcome better than the twenty more years that would be lived by Jane rather than by Sally.

Suppose next that we are considering this case at some much later time. Jack is aged seventy, and is about to die. We know that we could choose whether the outcome will be A or B. We could either enable Jack to live for one more year, or cause it to be Jane rather than Sally who will later exist and live for twenty more years. Since Jack is about to die, we may believe that we ought to save Jack's life, even though we would give him only one more year. On this view, our moral reason to save Jack's life would be stronger than our reason to cause it to be Jane rather than Sally who would later exist.

Though this belief is fairly plausible, we could not justifiably change our view about the relative goodness of these outcomes. This relative goodness cannot depend either on *when* we compare these outcomes, or on which of these outcomes would be what actually happens, or comes about. If it was earlier true that outcome A would be better than outcome B, it could not later be true that outcome B would be better than outcome A, nor could it be true that B would have been better than A. If we believe that, when Jack is about to die, we ought to save Jack's life, we should claim that we ought to save Jack's life even though this is not the act that would make things go best.<sup>19</sup>

We might then believe that

(P) we ought to give some benefits to actual people rather than giving some greater benefits to people who are now merely possible and who, if we did not give them these benefits, would never exist.

We might similarly believe that

(Q) we ought to give some benefits to presently existing people rather than giving some greater benefits to future people.

Though these beliefs are fairly plausible, they are about the rightness of acts, not the goodness of outcomes. We could not defensibly believe that, in giving such lesser benefits to actual people, or to presently existing people, we would be making things go better. If we gave similar benefits to people who are now merely possible, or to future people, these benefits

19. At this point in the text, Parfit had a note to himself and to readers of the draft that read "[some material to be added here]."

would do as much to make things go better. But we might defensibly believe that we ought to give such lesser benefits to presently existing people even though these acts would make things go worse. These beliefs would be like the view that we ought to save our children's lives rather than saving the lives of more children who are strangers to us. This view does not imply that it would be better if it would be our children's lives that were saved. We can believe that we ought to save our children's lives, though it would be worse if more children die. We are sometimes morally required to act in ways that would not make things go best.

Though we can plausibly believe that we ought to give such lesser benefits to presently existing people, rather than giving greater benefits to other people, we should remember that, if we and many others often act in such ways, we and others may together make things go much worse. We might save some present people from harm rather than saving future people from some greater harms. Our successors might do the same, and their successors might do the same, thereby making the quality of future people's lives lower, and lower, and lower. With such acts, we and our successors might together wreck the Earth.

#### VI. DIFFERENT NUMBER CASES

I have been discussing Mixed Same Number Cases. In such cases, though no one would exist in all of the outcomes that we are considering, some people would exist in more than one of these outcomes. We should also consider cases in which no one would exist in more than one outcome. In such cases, no outcome would be worse for anyone than any other outcome. When we think about the further future, most of the important questions are about such cases.

Suppose that, in *Case Six*, either

A: Many people will exist whose quality of life would be very high

or

B: Many other people will exist whose lives would be barely worth living.



Neither of these outcomes would be worse for anyone. Though the people in B would not have good lives, it would not have been better for them if they had never existed. Since B would be worse for no one, the Strong Narrow Principle implies that B could not be worse than A, and the Weak Principle fails to imply that B would be worse. It is clear, however, that B *would* be worse than A. This case provides a strong objection to these Narrow Principles. This objection is stronger if, as I have argued, we can reject these Narrow Principles but can still explain this part of morality—the part concerned with well-being—in person-affecting terms. We can claim that there can be intrinsic existential benefits, since it is good to be caused to exist and have a life that is worth living. Since the people in B have lives that are barely worth living, this outcome gives these people only weak existential benefits. Outcome A would give to as many other people great existential benefits. B would be worse than A, as the Wide Principle implies, because B would be much less good for people. Compared with existing and having some high quality of life, it would be much less good to exist and have a life that is barely worth living.<sup>20</sup>

Different Number Cases raise several other new questions. To introduce one of these questions, we can suppose that, in *Case Seven*, two possible outcomes are:

A: Tom will live to 40	Dick will live to 40
B: Tom will live to 80	_____

I believe that, of these outcomes, B would be better than A. That is what most of us would believe.

As one way to defend this belief, some people might appeal to the Two-Tier View. B would give Tom an ordinary comparative benefit of forty more years of life. A would give Dick a non-comparative existential benefit of forty years of life. On the Two-Tier View, comparative benefits have greater moral weight, by doing more to make outcomes better. If comparative benefits have even slightly more weight, B would be better than A.

20. Parfit here had a note to himself and to readers of the draft that read “[material to be added].”

This view, I have claimed, is false. Most benefits are in one way comparative, since we receive such benefits when the alternatives would have been worse for us. But these benefits are comparative only in the sense that they are the greater of two intrinsic benefits. These benefits are not essentially comparative. In *Case Seven*, there are two possible intrinsic benefits. Outcome A would give Dick the existential benefit of forty years of life, and outcome B would give Tom the comparative benefit of forty more years of life. This benefit to Tom would not be much greater, or do much more to make the outcome better, than this benefit to Dick. Dick's forty years of life would not be made to be less good for him by the fact that Dick's non-existence would not have been worse for him or bad for him. Since existential benefits are intrinsic, such benefits do as much as comparative benefits to make outcomes better.

There is a different way to explain why, in *Case Seven*, B would be better than A. In outcome A, Tom and Dick would together live for a total of eighty years. In outcome B, Dick would not exist and Tom would live for eighty years. Since A would benefit Dick about as much as B would benefit Tom, these outcomes would give both people roughly the same total sum of benefits. We can ask

Q4: Would it be in itself better if the same sum of benefits came to fewer people?

The rough answer, I believe, is Yes. We can claim that

(R) when two outcomes would give people the same total sum of benefits, it would be in one way better if these benefits were shared equally between fewer people.<sup>21</sup>

If Tom lives to forty and Dick lives to forty, that would be less good than if Tom lives to eighty and Dick never exists.

There are other, clearer cases. In considering Q4, we can imagine people whose lives could go better or worse, not merely by being longer or shorter, but in other ways. These people would have, in their lives as a whole, different levels or amounts of well-being. People at higher

21. Compare the *Consolidate Additional Benefits View* in Temkin's *Rethinking the Good*, p. 68.

levels would have greater sums of benefits, which we can represent with numbers. We might suppose, for example, that

Adam is at level 100, Bill is at level 50, and Charles is at level 1.

In being at level 1, Charles would be not far above the zero level at which lives are not worth living. Bill's level would be much higher. Compared with Charles's sum of benefits, Bill's sum would be about fifty times greater. Adam's sum of benefits would be about twice as great as Bill's. These differences between these levels of well-being would be much less precise than these numbers suggest.

We can now suppose that, in *Case Eight*, either

A: One person will exist at level 100

or

B: Ten other people will exist at level 10.

These outcomes would each give people the same total sum of benefits. On the view summed up by (R), A would be better than B. That is what most of us would believe.

In considering this view, we can distinguish between two senses or ways in which one of two outcomes might benefit people more. One outcome would

benefit people more in the *collective* sense if this outcome would *together benefit people more*,

and

benefit people more in the *individual* sense if this outcome would *benefit each person more*.

We can first consider outcomes in which there would be no inequality between different people. Everyone would have the same quality of life or level of well-being. Inequalities raise complications that are best considered elsewhere. That is one reason why the answer to Q4 is only roughly Yes.

When we compare outcomes in which the same number of people would exist, and there would be no inequality between people, these two ways of *benefiting people more* always coincide. If one of two outcomes would together give people a greater total sum of benefits, and each person would get an equal share of this greater sum, each person would get an equal benefit. In such cases, we need not distinguish between these two ways of benefiting people more.

When we consider Different Number Cases, however, we have reasons to draw this distinction. In such cases, the Wide Person-Affecting Principle can take two forms. According to

*The Wide Collective Principle:* One of two outcomes would be in one way better if this outcome would together benefit people more, by giving people a greater total sum of benefits.

According to

*The Wide Individual Principle:* One of two outcomes would be in one way better if this outcome would benefit each person more.

When I first discussed wide person-affecting principles, I considered only this Collective Principle.<sup>22</sup> This principle, I claimed, might solve the Non-Identity Problem. But I then rejected this principle. According to

*The Impersonal Total Principle:* It would always be better if there was a greater total sum of well-being, such as a greater sum of happiness.

This principle seems to imply

*The Repugnant Conclusion:* Compared with the existence of many people whose quality of life would be very high, there is some much larger number of people whose existence would be better, even though these people's lives would be barely worth living.

These very many people's lives might together contain a greater total sum of well-being, just as there might be a greater mass of milk in a

22. Parfit, *Reasons and Persons*, pp. 396–401.

vast heap of bottles that each contained only one drop. The Wide Collective Principle, I assumed, restates the Impersonal Total Principle in person-affecting terms. This person-affecting version of the Total Principle seemed to provide a stronger argument for the Repugnant Conclusion. Compared with denying that outcomes would be better if they increased the total sum of well-being, it seemed harder to deny that outcomes would be better if they benefited people more. That might be true, I assumed, if there existed enough people who each received the slight benefit of having a life that was barely worth living. Because the Repugnant Conclusion seemed to me very implausible, I claimed that we ought to reject this Wide Collective Principle.

This claim made two mistakes. We cannot justifiably reject strong arguments merely by claiming that their conclusions are implausible. And rather than rejecting this Collective Principle, I should have combined this principle with the Individual Principle. According to this

*Wide Dual Person-Affecting Principle:* One of two outcomes would be in one way better if this outcome would together benefit people more, and in another way better if this outcome would benefit each person more.

This principle is, I believe, the best version of the Wide Principle that I have been discussing. This principle appeals, I shall more briefly say, both to *the benefits to each* and to *the benefits to all*.

In some cases, one of two outcomes would be better in only one of these ways. This Dual Principle plausibly implies that

if  $N$  people existed at level 100, this would be better than if  $2N$  people existed instead at level 50.

The first outcome would be better because the benefits to each would be greater and the benefits to all would be the same. This principle also plausibly implies that

if  $2N$  people existed at 100, this would be better than if  $N$  people existed at 100.

The first outcome would be better because the benefits to all would be greater and the benefits to each would be the same.

In many other cases one of two outcomes would be better in one of these two ways, but worse in the other way. We can then ask which outcome would be better all things considered.

We can plausibly believe that slightly greater benefits to each would be outweighed by much greater benefits to all, and slightly greater benefits to all would be outweighed by much greater benefits to each.

Suppose next that, in *Case Nine*, either

A:  $N$  people will exist at 100

or

B:  $2N$  people will exist at 75.

The two parts of the Wide Dual Principle here more strongly conflict. Outcome A would benefit each person much more, but outcome B would together benefit people much more. We may find it harder to decide which of these outcomes would be better. To answer such questions, we must compare the relative importance of these two ways of benefiting people more. I can only start to do that here.

Suppose that, in *Case Ten*, either

A: One million people will exist at level 1,000

or

Z: One hundred billion people will exist at level 1.

As these numbers overly precisely imply, the lives of the people in outcome Z would not be well worth living, since these lives would not be much above the zero level at which lives are not worth living. The lives of the people in outcome A would be very well worth living. In each of these outcomes, the total sum of benefits would roughly correspond to the number of people who would exist multiplied by their level of well-being. In outcome A, this total sum of benefits would be roughly a thousand million, which is a billion. In outcome Z, this total sum would

be greater, since it would be roughly a hundred billion. If we appealed only to the Wide Collective Principle, we would have to conclude that Z would be better than A. This would be one version of the Repugnant Conclusion.

The Wide Dual Principle does not imply this conclusion. This principle appeals not only to collective but also to individual benefits. In *Case Ten*, Z would be collectively better than A, since Z would together benefit people roughly a hundred times more. But Z would be individually worse than A, since Z would benefit each person roughly a thousand times less. The lives of the people in A would be roughly a thousand times more worth living. We can plausibly believe that this second fact would do much more to make this outcome better. On this version of the Wide Dual Principle, Z would be worse than A.

Compared with the existence of many people whose quality of life would be very high, it would be worse if there existed instead a hundred thousand times as many people whose lives would be barely worth living. That is what most of us would believe. On this view, we might roughly say, though it matters whether what happens would together benefit people more, it matters more whether what happens would benefit each person more.

Here is another way to ask what we believe. According to

*The Repugnant Conclusion:* Compared with the existence of many people whose quality of life would be very high, there is some much larger number of people whose existence would be better, even though these people's lives would be barely worth living.

This conclusion seems repugnant because there seems to be little that is good in a world in which everyone's life would be so close to the zero level. Consider next this

*Analogous Conclusion:* Compared with the existence of many people whose lives would be barely worth living, there is some much higher quality of life whose being had by everyone would be better, even though the numbers of people who exist would be much smaller.

Few of us would find this conclusion repugnant. Many of us would believe that, if everyone's quality of life would be very high, this

outcome would be very good, though it would be even better if more such lives were lived. On this view, we would believe that in *Case Eleven*, compared with the existence of

one hundred billion people whose lives would be barely worth living,

it would be better if there existed

one million people whose quality of life would be roughly a thousand times higher.

This outcome would be better even though, because many fewer people would exist, the total sum of benefits would be roughly a hundred times smaller.

As these remarks show, the Wide Dual Principle does not imply the Repugnant Conclusion. This principle would often imply that one of two outcomes would be better though the people who exist would have a much smaller total sum of benefits. Such outcomes might be better, though these benefits would come to many fewer people, because these people's quality of life would be higher.

These remarks do not show that, if we can justifiably appeal to this Wide Principle, we can thereby avoid the Repugnant Conclusion. If we believe that it would always be in one way better if there existed more people who would together have a greater total sum of benefits, this belief might be one premise of more complicated and forceful arguments for the Repugnant Conclusion. But this Wide Principle would be only one of our beliefs. We might also justifiably believe that great losses in the quality of people's lives could not be outweighed by any increase in the sum of benefits, if these benefits came in the lives of people whose quality of life would be much lower. I have started to defend this belief elsewhere.