THE MORALITY OF CAUSING PEOPLE TO EXIST

Jeff McMahan

Most of us have firm moral beliefs about causing people to exist. We believe, for example, that we are not obligated to have children simply because those children could be expected to have lives worth living. Indeed, while we accept that people have a right to have children, we also believe that that right must be exercised in a way that respects the fact that there are strong moral reasons *not* to have children, or at least that there are reasons for each couple to limit themselves to one or at most two children. For we believe that further increases in the number of people on earth are in general undesirable, primarily because of the deleterious impact that expanding human populations have on the natural environment. Thus respect for the interests of those who already exist requires that there be some moral restrictions on procreation. In fact, most of us believe that procreation should be limited not just for the sake of those people who presently exist but also for the sake of future generations. We think that there is a moral reason to limit the size of the human population now in order to protect the environment for the sake of those people who will exist in the future.

The purpose of this paper is to indicate how extraordinarily difficult it is to provide a coherent defense of these and other common beliefs. To many people, this will be surprising. They will think that there is a simple and obvious way to defend the common sense view. This is to appeal to the principle that we are required to take account of a person's interests or welfare only if he is at some time *actual* - in particular, only if he now exists or will exist in the future. To many people, it seems virtually self-evident that, while we must respect the interests of both presently existing people and future people, merely possible people simply do not count. Thus, in deciding whether to bring a certain person into existence, we must take into consideration the impact that his existence would have on the interests of present and future people. But it is not a reason to cause him to exist that, if we were to do so, his own life could be expected to be worth living. For he is merely a possible person and need not be considered in our deliberations.

Many people have found these simple claims persuasive. But if we look deeper, problems emerge. How, for example, are we to understand the terms "future people" and "possible people"? It is frequently assumed that future people are simply those people who *will* exist. Whatever people will in fact exist in the future are now future people relative to us. The term "possible people" is also normally understood quite literally: these are people who might or might not exist in the future. But "might or might not" depending on what? The indeterminacy here cannot be merely epistemic. It is not just that we cannot now be certain whether or not they will exist. For that is true of future people as well. Rather, a person is normally classified as a possible person relative to some choice on which his existence depends. Let us say that the existence of a person P depends on choice C if there are some possible outcomes of C in which P would exist and some in which he would not. P is then a possible person relative to C.

Although it is common, this way of drawing the distinction between future and possible people is inadequate for the purpose of defending common sense moral beliefs. For, given these definitions, one cannot say that the interests of future people count while those of possible people do not. This is because the two categories overlap. If a person's existence depends on my choice, he is a possible person relative to that choice. But suppose that the outcome of my choice is that he does in fact come into existence. In that case, he was both a possible person *and* a future person at the time that I made my choice. For his existence depended on my choice but, given that he did later come to exist, it was also true at the earlier time (given certain plausible assumptions about the truth of future-tense propositions) that he *would* later exist. But if the same person can be both a possible person and a future person, then it is incoherent to claim that the interests of all future people must be taken into account while the possible interests of possible people are morally irrelevant. For this implies that the same person's interests must simultaneously be taken into account and not be taken into account.

Some writers have sought to evade this incoherence by assuming that a person's status as possible or future can shift, not just from one decision-making context to another (as it must if a person's status as a possible person is always relative to a particular choice) but from one time to another. Thus one writer claims that, "once we do decide to have children then we are obligated to begin taking their interests into account, and this process may even lead to a reversal of that decision."[1] The assumption here is that a child that one might cause to exist counts as a possible person until one decides to cause her to exist. At that point, she then becomes a future person whose interests must be taken into account. This, however, is obviously confused. Merely to decide to cause a person to exist does not make that person a future person - that is, it does not guarantee that that person *will* exist. This is shown by the fact that one can change one's mind. If a person's status as future or possible could shift in this way with a change of the agent's mind, then the decision-making process could continue forever in a circle: since a person that one might cause to exist is a possible person whose interests do not count, one can cause her to exist; but once one decides to cause her to exist, she becomes a future person whose interests do count; if her life would be bad, one ought not to cause her to exist; but if one then changes one's mind and decides not to cause her to exist, she becomes a possible person again; therefore her interests do not count and it again becomes permissible to cause her to exist; and so on forever.

Perhaps because of problems like these, some writers have drawn the distinction between future and possible people differently. While they have retained the definition of a possible person as someone whose future existence depends on the outcome of a particular choice, they have suggested that the classification of a person as a "future person" should also be relative to particular choices. The alternative understanding of the notion of a future person contains three elements. According to this understanding, a person P is a future person relative to choice C if [1] he does not now exist and [2] his future existence does not depend on the outcome of C. Some have thought that the third element is that [3] P would exist in *all* the possible outcomes of C - that is, that he *will* exist in the future. Then the relevant moral claim is, as Peter Singer puts it, that "a life does not count unless the person exists or will exist *independently of our action*."[2]

But whether a person *will* exist in the future is something that cannot be known and may not even be determined at the time that choice C is made. So, if the classification of a person as a future person is to serve as a guide to action, perhaps the third element in the definition should be that [3] there is a certain probability that P will exist in the future. Then the extent to which P's interests should count in the agent's deliberations about C might be discounted for the probability that P will not in fact exist. But the details need not detain us here. The important point is that P counts as a future person relative to C if he may exist in the future and whether or not he will exist is independent of the outcome of C. Given this understanding of the notion of a future person, we can now coherently assert the following principle:

> *The Independence Requirement*: In making a certain choice, one must take account of the interests of existing people and all those who are future people relative to that choice, though one should ignore the possible interests of those who are possible people relative to that choice.

This principle is coherent because the categories of future and possible people no longer overlap. If a person's existence depends on one's choice, he is a possible person relative to that choice; if there is some chance that he will exist and his existence is independent of one's choice, then he is a future person relative to that choice.

Despite the fact that the Independence Requirement is coherent and seemingly plausible, it is unacceptable, for a variety of reasons. First, as Derek Parfit has shown, there are surprisingly few future people.[3] What I mean by this paradoxical claim is that people who will exist in the future very often count as possible people relative to the choices that we make now. For who will exist in the future depends on a far greater range of acts than most of us have realized. This is because who will exist in the future is determined by which human gametes will be joined together and this in turn depends on a wide variety of circumstances. If, for example, my mother had decided to go to a different university from the one at which she met my father, I would never have existed. Thus I was a possible person relative to her choice of university. If, after his marriage, my father had taken a different job from the one he actually took, this would have affected my parents' lives in countless ways - in particular, it would have affected the time at which they conceived their first child. Their first child would have been conceived from different gametes and thus would not have been me but someone else instead. I was therefore a possible person relative to my father's choice of employment. And so on for an indefinite number of other choices that my parents and others made prior to my conception.

For the same reasons, choices between social policies contribute to determining who will exist in the future. If a society chooses to adopt Policy 1 rather than Policy 2, this will have widespread effects on the details of people's lives. People will engage in different activities from those they would have engaged in had Policy 2 been adopted instead. They will meet different people, marry different people, and, even within those marriages that would have been the same, people will conceive their children at different times. These effects increase over time. After a certain time has passed, virtually all of those who will exist in the society given that Policy 1 was chosen will be different from

those who would have existed at that time had Policy 2 been chosen instead. Hence virtually all of the people who exist at that time will have been possible people relative to the choice of Policy 1.

It is important to note that Policy 1 and Policy 2 could be policies that would have very different effects on the environment or on population growth. Indeed, they could be competing environmental policies, or population policies. They could therefore have dramatically different effects on the quality of human life in the future. But each would lead to the existence of different people in the future. Those people who will exist only if Policy 1 is adopted and those who will exist only if Policy 2 is adopted are all possible people relative to the choice between 1 and 2. According to the Independence Requirement, therefore, these possible people should be ignored in deliberations about the choice between 1 and 2. This implication can be generalized. In our deliberations about matters of social policy, the Independence Requirement implies that we can ignore the effects of any policy on those members of future generations whose existence depends on whether or not the policy is adopted. In many cases, this will include most of the people who would exist in the future if the policy were adopted. In short, the Independence Requirement implies that our concern for future generations should be much more limited in scope than in fact it is. If we believe that the well-being of distant future generations matters, then we must reject the Independence Requirement.

The foregoing argument may be difficult to understand. I will explain it in a different way, which I hope will make it clearer. Suppose that we can choose to do one of two acts, A1 and A2. If we do A1, this will cause person P1 to exist. If we do A2, this will cause person P2 to exist. P1 and P2 are possible people relative to the choice between A1 and A2. Therefore, according to the Independence Requirement, what their respective lives would be like is morally irrelevant, except insofar as their lives would affect those of existing people or those who are future people relative to this choice. This, however, seems quite implausible. Suppose that P1 would have a much better life than P2, perhaps because his or her physical or cognitive capacities would be higher or simply because the circumstances of his or her life would be better. This is surely a relevant consideration. If all other things are equal, it seems that it would be wrong to cause P2 to exist rather than P1. In general, if there is a choice between causing a happier person to exist and causing a less happy person to exist and other things are equal, it is wrong to cause the less happy person to exist.

A similar claim holds for groups. Suppose that, if one does act A1, this will cause a certain group of people, G1, to exist. If one does A2 instead, this will cause a different group of people, G2, to exist. Suppose that [1] G1 and G2 would contain the same number of people, [2] none of the members of G1 would also be in G2, [3] the members of G1 would be better off than the members of G2, and [4] other things would be equal.[4] In these conditions, it would be wrong to do A2, causing G2 to exist rather than G1. Yet, because the members of G1 and G2 would all be possible people relative to the choice between A1 and A2, the Independence Requirement implies that there is no reason to choose A1 rather than A2. This is one reason why the Independence Requirement is implausible. This objection is, moreover, a general statement of my earlier claim that the Independence Requirement is quite liberal in permitting us to ignore the well-being of

future generations. For A1 and A2 might be alternative environmental or population policies and G1 and G2 might be alternative future populations.

That the Independence Requirement allows us to cause less happy rather than happier people to exist, and thus allows us considerable latitude in choices among social policies when those policies will affect the composition of future generations, is one intuitive objection to the Requirement. But there are others as well. For example, suppose that it was certain that, if one were to conceive a child, that child would have Lesch-Nyhan Syndrome, or perhaps a severe form of Dystrophic epidermolysis Bullosa affecting not just the skin but the linings of the digestive and respiratory tracts as well.[5] In either case, the child's life would be extremely short and filled with terrible pain that would not be offset or compensated for by any of the ordinary goods of infant life. These are conditions, in short, that virtually guarantee that life cannot be worth living. Most of us believe that, if it were certain that any child one might have would have one of these conditions, then it would be wrong to have children. This is because we believe that it is wrong to cause someone to exist knowing that that individual's life would not be worth living - indeed, would be "worth *not* living," or worse than no life at all. Yet such a person would be a possible person relative to one's choice between causing him to exist and not causing him to exist. The Independence Requirement therefore holds that, unless his existence would be worse for existing or future people, there is no moral reason not to cause him to exist. This is clearly unacceptable.

There is an important general lesson to be drawn from these intuitive objections to the Independence Requirement. This is that neither of the two proposed ways of distinguishing between future and possible people provides a basis for the original claim that, in deliberating morally about what to do, one should consider only those people who exist at some time - in particular, only those people who now exist or definitely will exist. For, given the way that possible people were defined, some possible people will in fact exist. In other words, some people who *might* exist in the future *do* in fact later exist. Therefore a principle, such as the Independence Requirement, that says that one should ignore all possible people is in effect saying that one should ignore some people who will in fact exist. And this is incompatible with the original claim that all those who *will* exist must be taken into account.

We cannot, however, simply start over by reverting to the original claim that we must take account of *all and only* those who do or will exist. This is not a principle that is capable of guiding one's action. For some choices one makes determine who will exist. In making such a choice, one cannot be guided by a principle that says that one must take account only of those who do or will exist. For who will exist is something that will be determined by the outcome of one's choice and therefore cannot be included among the data on which the choice is to be based.

Given that it is not possible to guide one's action by the principle that one must consider only those who do or will exist, it is not surprising that people shifted, perhaps without being aware of it, to the different view that one must not, in making a certain choice, consider those whose existence depends on that choice. For this latter view - namely, the Independence Requirement - has a certain *prima facie* plausibility. If

someone's existence depends on one's choice between acts, then it can be argued that one's choice cannot be worse for that person, even if one's choice in fact results in the person's existence. If an act cannot be worse for those who are possible people relative to the choice between doing it and not doing it, then there can be no reason not to do it *for the sake of those possible people*. If one chooses to do the act, it cannot be wrong because of its effects on those who are possible people relative to one's choice to do it.

While it may seem that these claims support the Independence Requirement, they are in fact part of the objection to it. To see why, we need to understand *why* an act cannot be worse for those who are possible people relative to the choice to do it. Then we need to see what this implies.

Consider first cases in which an act causes a person to exist with a life that is worth living. That person is a possible person relative to the choice to do that act. The reason why this act cannot be worse for the person is that, if the act had not been done, *he* would not have existed at all. If his life is worth living and the only alternative was for him not to exist at all, then the act that caused him to exist cannot have been worse for him. For it cannot be worse to exist with a life worth living than not to exist at all.

Acts of this sort can, however, be clearly morally wrong, and for reasons connected with their effects on possible people. For example, it can be wrong, as we have seen, to cause a less happy person to exist when it would have been possible to cause a much happier person to exist instead, even if the less happy person has a life that is worth living.[6] We cannot explain why such an act is wrong without taking into account its effects on possible people.

Next consider acts that cause people to exist whose lives are *not* worth living. It is true that an act of this sort cannot be wrong because it is worse for the person who is caused to exist. This seems to be true as a matter of logic. Even if a person's life would not be worth living, it is *necessarily* true that causing him to exist would not be *worse* for him. For the alternative to causing him to exist could not possibly be better *for him*, since in that alternative he would not exist at all. Since it is logically true that an outcome can be *worse* for someone only if there is an alternative that would be *better* for him, it follows that causing someone to exist cannot be worse for him, even if his life would not be worth living.

This seems to be true *only* as a matter of logic. As a substantive claim, it is misleading. For to say that causing such a person to exist cannot be *worse* for him suggests that it cannot be *bad* for him. But this does not follow. The claim that an outcome is worse for someone necessarily involves a comparison with other outcomes. But the claim that an outcome is bad for someone does not necessarily involve a comparison with an alternative outcome. Thus, even if it is true that it cannot be *worse* for a person to be caused to exist with a miserable life (since the alternative would not be better for him, for in that alternative he would never exist at all), that is compatible with the claim that this could be *bad* for him (since it would cause him to have a miserable life).

I believe that it *is* bad for a person to be caused to exist if his life is miserable, or not worth living. And I believe, as I think most people do, that it can be wrong to cause such a person to exist. If this is right, then possible people *do* matter morally. They do count. If one refrains from causing a person to exist on the ground that, if he were to exist, his life would be miserable, then one's action is being guided by respect for the possible interests of a merely possible person - one who will in fact will never exist.

This, I think, is common sense. Most of us think that we must take account of possible people in our moral deliberations. While most people think that there is no moral reason to cause a person to exist just because his life would be worth living, this is not because that person is merely a possible person. For in other cases - for example, causing a person to exist with a life that would not be worth living - we accept that possible people matter. The common sense view is:

> *The Asymmetry*: While the expectation that a person's life would be miserable, or not worth living, provides a strong moral reason not to cause him to exist, the expectation that a person's life would be happy, or worth living, provides *no* moral reason to cause him to exist.

One cannot defend the second of the two claims that comprise the Asymmetry - that is, the claim that there is no reason to cause a happy person to exist for his own sake - by arguing that this person does not now exist, or that he is merely a possible person, or that it would not be worse for him if one were not to cause him to exist. For parallel points are also true of the first claim. The person whose life would not be worth living also does not now exist; he is also a possible person; and (as noted above) it would not be worse for him if one were to cause him to exist. (And, of course, it would also not be better for him if one were not to cause him to exist.)

In short, we cannot hope to defend our intuitions by appealing to the distinction between future and possible people or to the view that the only people who count are those who are at some time actual. Sometimes possible people do count, even those who in fact never exist.

I suspect that the Asymmetry cannot in fact be defended. There is a simple argument, which I call the "Symmetry Argument," that challenges the Asymmetry. The Symmetry Argument purports to show that there *is* a moral reason to cause a person to exist if his life could be expected to be worth living. It begins with the first of the two claims that make up the Asymmetry - namely, the claim that it is wrong to cause someone to exist if his life would be miserable, or not worth living. Next we must ask why this is true. The most natural explanation appeals to the idea that to cause such a person to exist is to harm him. If it is *bad* for him to exist with a miserable life, then to cause him to exist must be to harm him. Next we appeal to the principle of nonmaleficence: other things being equal, there is always a moral reason not to do what harms people. The first part of the Symmetry Argument, then is this:

[1] To cause someone to exist with a miserable life is on balance to harm that person.

[2] There is a moral reason not to do what harms people, other things being equal.

[3] If one were to cause a certain person to exist, his life would be miserable.

[4] Other things are equal - that is, there are no countervailing reasons that outweigh or override the reason not to harm this person by causing him to exist.

[5] Therefore one ought not to cause this person to exist.

The next step in the Symmetry Argument consists in the claim that, if to be caused to exist with a life that is not worth living can be a harm, then it must also be true that to be caused to exist with a life that is worth living can be a benefit. To deny this - to accept that to be caused to exist can be a harm but to deny that it can be a benefit - would, in the absence of some plausible explanation, be unacceptably *ad hoc*. And it is difficult to imagine what sort of explanation could be given for the claim that there is an asymmetry of this sort. The moral implications of supposing that there is such an asymmetry are, moreover, extremely implausible. For example, if it is possible to harm people by causing them to exist but not possible to benefit them, then any decision to have a child will carry a risk of harming the child but will not involve a possibility of benefiting him. Thus, given these assumptions, there is always a presumption against having a child that has to be rebutted. Unless the expected benefits to existing people of having the child outweigh the risk of harm to the child, it is wrong to have children. But it seems clear that there is no such general moral presumption against having children. Thus it seems safe to assume that, *if* causing a person to exist can harm that person (if his life is not worth living), then it can also benefit him (if his life *is* worth living).

The next step in the Symmetry Argument is to note that, just as there is a general principle of nonmaleficence, so there is also a corresponding general principle of beneficence. This principle holds that, other things being equal, there is always a moral reason to do what would benefit people. We can now combine these claims to get the rest of the argument.

[6] To cause someone to exist with a life that is worth living is on balance to benefit that person. (As I noted, it seems that this must be true if [1] is true.)

[7] There is a moral reason to do what benefits people, other things being equal.

[8] If one were to cause a certain person to exist, his life would be happy, or worth living.

[9] Other things are equal - that is, there are no countervailing reasons that outweigh or override the reason to benefit this person by causing him to exist.

[10] Therefore one ought to cause this person to exist.

The last of these claims is generalizable. The reasoning in the Symmetry Argument leads to the general conclusion that there is always a moral reason to cause a new person to exist if his life would be happy, or worth living. It does not follow, of course, that this

reason is always decisive, so that one is obligated to cause new people to exist whenever this is possible. The Symmetry Argument allows that the reason to cause people to exist can be overridden by countervailing reasons. But, when it is possible to cause a person to exist who would have a life worth living and there are no sufficiently weighty countervailing reasons, either moral or prudential, then one ought to cause the person to exist.

Because the Symmetry Argument is incompatible with the Asymmetry, most people will want to reject it. And there are various ways in which one might try to refute it. One way to try to resist the Symmetry Argument is to appeal to the distinction between *doing* and *not doing*. It is often argued that we are morally more responsible for what happens or fails to happen as a result of what we *do* than we are for what happens or fails to happen as a result of what we do *not do*. If it can be either good or bad for a person to come into existence (depending on the quality or character of his life), then the coming-into-existence of a miserable person and the not-coming-into-existence of a happy person are both bad outcomes. Bringing a miserable person into existence is a case of *doing*, while not bringing a happy person into existence is a case of *not doing*. Thus, assuming that we are more responsible for the results of what we do than we are for the results of what we do not do, it follows that to bring a miserable person into existence is worse, other things being equal, than not to bring a happy person into existence.

The problem with this strategy for resisting the Symmetry Argument is that it requires an unacceptably strong form of the claim that there is a moral asymmetry between doing and not doing. In order to deny that it is wrong, other things being equal, to fail to cause a person to exist whose life would be worth living, one would have to hold that we cannot be held responsible for what happens or fails to happen as a result of our not doing something. Even if it is qualified to allow that we can have special positive obligations, such as those derived from promising, this denial of responsibility for the results of what we do not do is too strong. It implies, for example, that, in the absence of a special obligation, there is no moral reason to prevent other people from suffering harm. But, if we accept that there are moral reasons to do things for people that do not arise from special obligations, then we cannot evade the Symmetry Argument by appealing to the distinction between doing and not doing.

A second way of trying to resist the Symmetry Argument is to appeal to a theory of rights. Perhaps one can say that, if one causes a person to exist with a life that is not worth living, one will have violated that person's rights.[7] But it cannot be plausibly claimed that, if one fails to cause a person to exist when that person would have had a life worth living, one will have violated that person's rights. For those who never exist cannot have rights. But, while the appeal to rights thus excludes the possibility that one could have a duty, based on a respect for people's rights, to cause a person to exist whose life would be worth living, it does not exclude the possibility that there might be some *other* moral reason to bring him into existence. To support the conclusion that it cannot be wrong, other things being equal, not to bring such a person into existence, one would need the further claim that an act cannot be wrong unless it violates a right, and this seems false. Not all wrong acts violate rights.

Let us now consider a third way of resisting the Symmetry Argument. This is to deny that there is a general principle of beneficence that corresponds to the general principle of nonmaleficence. In other words, one might argue that there is a general moral reason not to do what would harm people and yet deny that there is a general moral reason to benefit people. This view once found expression in theories that were labeled "negative utilitarianism."

Again, however, to support the view that there is *no* moral reason to cause a person to exist on the ground that his life would be worth living, this view must take an unacceptably strong form. It must hold that there is never any moral reason (except perhaps in cases involving special obligations) to do what benefits people. Again, this seems false.

It may be, however, that a weaker version of this view is true. Most of us believe that it is in general more important to prevent, to alleviate, or not to cause misery or unhappiness than it is to provide or to preserve benefits. Of course, the distinction between benefiting and harming requires careful elucidation and should not be conflated with the distinction between doing and allowing.[8] For example, the relevant notions of benefit and harm have to be sensitive to people's overall levels of well-being, for the main intuition that this view must capture is that those who are in general badly off have priority over those who are in general well off. But, if we could formulate a view that would capture this and other related intuitions, it would presumably imply that it is in general more important to prevent people from being miserable than it is to make people happy. And this would give us a weak version of the Asymmetry:

> *The Weak Asymmetry*: While the expectation that a person's life would be miserable, or not worth living, provides a *strong* moral reason not to cause him to exist, the expectation that a person's life would be happy, or worth living, provides only a *weak* reason to cause him to exist.

The main justification for the claim that the moral reason to cause a happy person to exist is weaker than the reason not to cause a miserable person to exist is that the prevention of misery matters more than the promotion of happiness. But there could be additional elements to the justification. If, for example, it is true that we are more responsible for the results of what we do than we are for the results of what we do not do, this could be a reinforcing part of the explanation why it is worse to cause a miserable person to exist than it is not to cause a happy person to exist. And one might also invoke Kant's distinction between perfect and imperfect duties. One could claim that, while the duty of nonmaleficence is a perfect duty, the duty of beneficence is an imperfect duty. According to this view, the duty to benefit people does not apply at all times - not even on all occasions when there are no countervailing reasons. While there is in general a duty to bestow benefits, one is allowed some discretion in the choice of both the occasions and the beneficiaries. And there is a limit to the benefits that one is required to bestow.

Many people will object to the Weak Asymmetry. One reason for objecting to it is that it fails to discriminate between benefiting someone by causing him to exist and

benefiting existing people. If the amount of benefit would be the same, the Weak Asymmetry seems to treat these two ways of benefiting people as morally equivalent. Yet even those who believe that there is a reason to benefit people by causing them to exist also tend to believe that it is more important to bestow benefits on existing people than it is to benefit new people by causing them to exist.

It is tempting to try to defend the common view by reintroducing the idea that possible people matter less. Rather than claiming that possible people do not count at all, perhaps one could claim that the interests of possible people count for less than those of existing people. *How much less* their interests count would be a matter for dispute. Suppose - to adopt a simple solution - that the interests of possible people count half as much as the interests of actual people. In that case, given a choice between benefiting an existing person and benefiting a possible person by causing him to exist, the benefit to the possible person would have to be more than twice as great as the benefit to the existing person in order for it to be preferable to benefit the possible person.

It might seem that this proposal for discounting the interests of possible people would make the Weak Asymmetry more palatable. But this is an illusion. For it is not acceptable to claim, as a general rule, that the interests of possible people count for less than the interests of existing people. This becomes clear when we consider the possibility of causing miserable people to exist. Again, suppose that the interests of possible people count only half as much as the interests of existing people. And suppose further that I have a choice between two options, both of which involve the prevention of suffering. In one option, I can (without violating any actual person's rights) prevent the existence of a possible person whose life would have an overall net value of negative 100.[9] Imagine that this would be a child of the sort mentioned earlier - that is, a child with a condition involving severe congenital mental retardation and chronic severe pain that could be fully relieved. The child would live for only a couple of years, years filled with suffering uncompensated for by any of the things that normally make a life worth living. In the other option, I can prevent an existing person from experiencing a significant amount of suffering. This suffering would reduce his overall well-being by 55. Obviously, if we think that the existing person's interests count for twice as much as those of the possible person, then I should prevent the suffering of the existing person rather than preventing the existence of the child, whose suffering would be nearly but not quite twice as great. Intuitively, this seems to be the wrong choice.[10] In this case, the interests of the possible person seem to matter just as much as the interests of the existing person.

In spite of this, most of us continue to feel that it is more important to benefit existing people than it is to bestow equal benefits on possible people by causing them to exist. An alternative way of explaining and justifying this belief is to appeal to the idea that special relations justify some degree of partiality. Most of us believe that morality allows us to assign greater weight to the interests of those to whom we are specially related than to the interests of strangers. So, for example, I am permitted to give some degree of priority to the interests of my family and friends simply because they are *my* family and friends. And it might be argued that I am similarly more closely related to people who exist now than to someone who does not yet exist at all. If so, then this alone

may justify the belief that it is more important to benefit existing people than it is to benefit possible people by causing them to exist.

There are at least two objections to this. Though each is significant, neither seems to me to be decisive. First, if the relation that we bear to those who exist but not to those who do not now exist is sufficiently important to justify our attributing greater weight to *benefiting* existing people, does it also make it more important not to *harm* existing people than not to harm possible people by causing them to exist? If it does, then this reduces the plausibility of the appeal to partiality, since (as we have seen) most of us believe that it is just as important to avoid harming people by causing them to exist as it is to avoid harming existing people.

The reason that this objection is not decisive is that, *even among existing people*, special relations do not tend to justify partiality in cases involving causing harm. They tend to justify partiality only in the distribution of benefits. Thus, while one is often entitled to use one's resources to favor one's friends and relations rather than using them to benefit strangers, it is in general no less wrong to inflict harm on a stranger than it is to inflict a comparable harm on someone to whom one is specially related. So it seems that the following two claims are consistent: [1] Partiality to existing people may make it more important to benefit them than to bestow comparable benefits on new people by causing them to exist. [2] It is just as bad to harm someone by causing him to exist with a miserable life as it is to inflict a comparable harm on someone who already exists.

A second possible objection to the appeal to special relations is that in some cases it may be true that the relation one bears to a possible person is more significant morally than the relation one bears to most existing people. This might be true, for example, if the possible person were one's possible child. It is, however, not clear that this really makes sense. If one's child does not yet exist, then one cannot *now* stand in a significant relation to that child, for there is no one to whom one could now be related in any way.

Because these two objections are not decisive, I am inclined to accept that it is justifiable to give some priority to benefiting existing people over causing happy people to exist. The justification appeals to the fact that one is related to existing people in a way that one is not related to possible people. If this is right, it may help to make the Weak Asymmetry more plausible. According to the Weak Asymmetry, there is a moral reason to cause a person to exist if his life would be worth living. But this reason is much weaker than the reason not to harm an existing person or the reason not to harm a person by causing him to exist with a miserable life. For, in general, it is less important to benefit people than it is not to harm them. And there is more reason to benefit those to whom we are specially related than there is to benefit those to whom we are not specially related. And we are not specially related to possible people at all.

I myself am inclined to think that the Weak Asymmetry is true. It is weak enough that, though it says that there is a moral reason to cause people to exist if their lives would be worth living, it does not commit us to the view that we should cause happy people to exist whenever possible. It is not, however, so weak that it is unproblematic. Rather, it leads to a serious problem. If we can benefit people by causing them to exist,

then surely there are circumstances in which the benefits to new people in causing them to exist can outweigh lesser losses to existing people. If the process of adding new people at some cost in the quality of life for existing people were repeated often enough, then the size of the world's population would increase while the average quality of life would decline. And this could be required by morality, assuming that there is a reason to cause new people to exist just because their lives would be worth living. Yet this is a conclusion that most of us are deeply reluctant to accept. We think that it is in general wrong to increase the size of the world's population if doing so would lower the quality of life for existing people.

There are in fact powerful arguments for the conclusion that we should increase the world's population that do not depend on the assumption that there is a moral reason to cause people to exist if their lives would be worth living. I will not rehearse those arguments here, for they are well known in contemporary moral philosophy.[11] My purpose in this paper has been to show how difficult it is to defend even the most basic common sense intuitions about the morality of causing people to exist. The reader should be aware that the problems I have considered are just the beginning.

NOTES

[1] Mary Anne Warren, "Do Potential People Have Moral Rights?", in R.I. Sikora and Brian Barry, eds., *Obligations to Future Generations* (Philadelphia: Temple University Press, 1978), p. 25. A similar confusion can be found in David Heyd, *Genethics: Moral Issues in the Creation of People* (Berkeley: University of California Press, 1992), pp. 100-103 and 129.

[2] Peter Singer, "A Utilitarian Population Principle," in Michael D. Bayles, ed., *Ethics and Population* (Cambridge, MA: Schenkman, 1976), p. 93. Emphasis added. By "our action" Singer of course means some particular act, not human action in general.

[3] Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), chapter 16.

[4] Compare the principle labeled *Q* in *Reasons and Persons*, p. 360.

[5] The appalling effects of the latter condition are described in Jonathan Glover, "Future People, Disability, and Screening," in Peter Laslett and James Fishkin, eds., *Justice between Age Groups and Generations* (New Haven: Yale University Press, 1992), p. 129.

[6] Another type of case is that in which an act is a necessary part of the cause of a person's existence but also has serious bad effects within that person's life. For examples of this kind of case, see Parfit's *Reasons and Persons*, section 126.

[7] I discuss the application of theories of rights to the problem of causing people to exist in "Problems of Population Theory," *Ethics* 92 (1981), pp. 124-27.

[8] For a careful analysis of the distinction between harm and benefit, see Joel Feinberg, *Harm to Others* (New York: Oxford University Press, 1984), especially chapter 4.

[9] The numbers are for purposes of illustration only and are not to be taken seriously.

[10] This argument reinforces Derek Parfit's arguments for what he calls the "No Difference View." See *Reasons and Persons*, pp. 366-71.

[11] These arguments are developed in Parfit's *Reasons and Persons*, chapters 17 - 19.