"Wrongful Life: Paradoxes in the Morality of Causing People to Exist" by Jeff McMahan, pp 445-475, in Bioethics /edited by John Harris, 2001, reproduced by permission of Oxford University Press.

14

WRONGFUL LIFE: PARADOXES IN THE MORALITY OF CAUSING PEOPLE TO EXIST

JEFF McMAHAN

I HARM AND IDENTITY

The issue I will discuss can best be introduced by comparing two cases, both of which involve a *Negligent Physician*.

The Preconception Case
A man and a woman are considering having a child but suspect that one of them may be the carrier of a genetic defect that causes moderately severe cognitive impairment. They therefore seek to be screened for the defect. The physician who performs the screening is negligent, however, and assures the couple that there is no risk when in fact the man is a carrier of the defect. As a result, the woman conceives a child with moderately severe cognitive impairments.

Had the screening been performed properly, a single sperm from the man would have been isolated and genetically altered to correct the defect. The altered sperm would then have been combined *in vitro* with an egg drawn from the woman and the resulting zygote would have been implanted in the woman's womb, with the consequence that she would later have given birth to a normal child.

Notice, however, that the probability is vanishingly small that the sperm that would have been isolated and altered would have been the very same sperm that in fact fertilized the egg during natural conception. And let us suppose that the egg that would have been extracted for *in vitro* fertilization would also have been different from the one that was fertilized during natural conception. In that case the child who would have been conceived had the screening been done properly would have developed from a wholly different pair of gametes and would thus in fact (even if it is not a matter of metaphysical necessity) have been a *different* child from the retarded child who now exists.

Now compare the Preconception Case with:

This chapter is an abridged and slightly revised version of a paper with the same title from Jules Coleman and Christopher Morris (eds.), *Rational Commitment and Social Justice: Essays for Gregory Kavka* (Cambridge: Cambridge University Press, 1998), 208–47.

The Prenatal Case
A physician negligently prescribes a powerful drug for a woman who is in the eighth month of pregnancy. The drug causes damage to the fetus's brain and the child to whom she gives birth is, as a consequence, moderately cognitively impaired.

Assume that each of us began to exist prior to eight months after the conception of his or her organism. (I have elsewhere defended at length the claim that each of us began to exist when the brain of his or her fetal organism developed the capacity for consciousness and mental activity—which happened sometime between the 20th and the 28th week of fetal gestation.)[1] In that case, it seems clear that the retarded child who develops as a result of the brain damage is numerically the same child as the child who would have existed had the damage not been done.

There is a sense in which the outcomes of the two cases are the same: a retarded child exists rather than a cognitively normal child. The difference is that, in the Prenatal Case, the retarded child and the possible normal child are the same child in two different possible lives, whereas in the Preconception Case the retarded child and the hypothetical normal child are different children. This latter fact poses a problem, which Derek Parfit calls the *Non-Identity Problem*.[2] Assume that, in both cases, the retarded child's life, though drastically limited in the goods it can contain, is not so bad as not to be worth living. If that is so, it seems that the Negligent Physician's action in the Preconception Case was not worse for the child. For if the Negligent Physician had not acted negligently, *that* child would never have existed; and to exist with a life that is worth living cannot be *worse* than never to exist at all.

So, in the Prenatal Case, the Negligent Physician's action was worse for the retarded child (because it caused him to be retarded rather than normal), while in the Preconception Case it was not. Is this a morally significant difference? And, if the objection to the Negligent Physician's action in the Preconception Case is not that it harmed or was worse for the child, what is it? How can we explain our sense that the Negligent Physician's action was morally objectionable in that case?

It is clear that, in both cases, the Negligent Physician's negligence has harmed the couple, who have been denied many of the joys of parenthood and who instead have the often anguishing burden of caring for a relatively unresponsive and highly dependent child. In the law they would be warranted in bringing a 'wrongful birth' suit against the physician, in which they

[1] See Jeff McMahan, *The Ethics of Killing* (New York: Oxford University Press, forthcoming 2001), chs. I and IV.
[2] Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), ch. 16.

as plaintiffs would claim damages for the harms his negligence has caused them.

But there seems to be more to it than this, even in the Preconception Case. Most of us believe that, quite independently of the impact of the physician's action on the parents, the retarded child in the Preconception Case ought not to have been caused to exist and that, given that he *has* been wrongfully caused to exist, the physician should be required to pay damages not only to compensate the parents for the injury done to them but also, insofar as possible, to enhance the life of the child. This latter conviction lingers even when we understand that, because of the Non-Identity Problem, the Negligent Physician's action was not worse for the child.

In discussing a similar pair of cases, Parfit claims that, given that the outcome is the same in each (that an individual is caused to exist with a disability), it makes no difference that in one case the outcome is worse for the individual whereas in the other case it is not. He calls this the *No-Difference View*. This view is plausible in our two cases. It implies that the objection to the Negligent Physician's action in the Preconception Case is equally strong as the objection to his action in the Prenatal Case. 'This suggests,' Parfit contends, 'that there is the same objection to each act.'[3] Since the objection to the physician's action in the Preconception case cannot be that it harmed or was worse for the retarded child, it follows, if Parfit is right, that this cannot be the objection to his action in the Prenatal Case either. Generalizing, Parfit claims that the fact that an effect is *worse for people*, or bad *for them*, is never part of the fundamental explanation of why the effect is bad. The area of morality 'concerned with beneficence and human well-being,' he writes, 'cannot be explained in person-affecting terms.'[4] It must instead be explained in *impersonal* terms.

The latter inference has broad-ranging implications for moral theory. I will explore these in some detail in Section III; for the moment I will illustrate the significance of Parfit's claim by noting one implication that is of particular interest. I cited earlier the view that each of us began to exist only when the brain of his or her physical organism developed the capacity to generate consciousness and mental activity. This understanding of personal identity, if correct, provides the basis for what seems to be a plausible argument for the permissibility of early abortion. For, according to this view, an early abortion does not kill one of us but instead merely prevents one of us from coming into existence. The organism that is killed is not numerically identical with the later person and thus is not deprived of the later life that is

[3] Derek Parfit, 'Comments,' *Ethics*, 96 (1986), 858.
[4] *Reasons and Persons*, 370–1.

precluded. Early abortion, then, is morally comparable to contraception: there need be no one for whom it is worse. The power of this view is illustrated by its ability to explain a common but otherwise puzzling judgment: namely, that it is less objectionable to *kill* a perfectly healthy early fetus than it is to injure or damage it in a comparatively minor way, e.g., a way that causes the subsequent person to have a minor physical disability. The explanation is that, provided that the abortion is desired by the parents, killing the fetus is not worse for anyone, while damaging the fetus harms the person to whose existence the fetus subsequently gives rise. In short, the instance of prenatal injury has a victim while early abortion does not.

According to the No-Difference View, however, that an act is bad or worse *for someone* is no part of the explanation why its effects are bad; accordingly, an act may have a bad effect, and thus be seriously morally objectionable, even if there is no one for whom it is worse or bad in any way. Hence, if Parfit is right, the fact that prenatal injury is worse for the future child does not explain why it is bad to injure a fetus; nor can one infer from the claim that an early abortion is worse for no one that it is not bad. For it might be bad impersonally. The No-Difference View thus appears to undermine both a seemingly plausible basis for distinguishing between prenatal injury and abortion and a powerful argument for the permissibility of early abortion.

II APPROACHES THAT IDENTIFY A VICTIM

Some writers have sought to address the threats that the Non-Identity Problem poses to common sense beliefs by arguing that, in such cases as the Preconception Case, a child born with a disability is adversely affected by the act that causes the disability, even if its life is worth living and it would never have existed had the act not been done. Imagine, for example, that yet another Negligent Physician gives a woman who is having trouble conceiving a child an inadequately tested fertility drug that both allows her to conceive a child and causes the child to suffer from some dreadful disease later in life. It seems reasonable to say that the Physician's act was the cause of the child's contracting the disease and that, in causing the terrible disease, the act *harmed* the child—even though it was not, on balance, worse for the child that the act was done. And, according to one proponent of this approach, 'that an agent is morally accountable for someone's suffering a harm, by virtue of having performed a certain action, seems a perfectly intelligible "person-affecting" explanation why his action is objectionable.'[5]

[5] Matthew Hanser, 'Harming Future People,' *Philosophy and Public Affairs*, 19 (1990), 59.

This approach has to be extended somewhat if it is to be applied to cases such as the Preconception Case. For, while contracting a disease is a discrete event that involves the worsening of a prior condition, congenital mental retardation is an inherent, constitutive aspect of a person's nature, not a contingent addition to his life. It is less easy, therefore, to regard the fact that the child in the Preconception Case is retarded as *a harm*. Still, we may invoke Joel Feinberg's notion of a 'harmed condition' in order to assimilate this case into the paradigm to which the vocabulary of harm is applicable. A harmed condition is 'a condition that has adverse effects on [an individual's] whole network of interests,' and is 'the product of a prior act of harming.'[6] Congenital retardation seems to doom many of the retarded individual's interests to frustration. And we may, if only for the sake of argument, grant that the Negligent Physician's action in the Preconception Case counts as an act of harming. The objection, then, to the Physician's action is that it causes the child to exist in a harmed condition. The child is therefore appropriately seen as the victim of this action.

It is not clear whether this approach, which may be called the *harm-based approach*, presupposes that it can be a harm to be *caused to exist* in a harmed condition. Some proponents of the harm-based approach might wish to avoid being committed to the claim that to be caused to exist can be either good or bad for a person. They will therefore want to claim that the Negligent Physician harms the retarded child not by contributing to causing his existence but instead by causing him to be cognitively disabled. It is not obvious, however, that this distinction is tenable, since *this* child can exist only if he is disabled and the actual effect of the Negligent Physician's action is simply to cause this child to exist rather than another child. But I will put this problem aside and assume for the sake of argument that the Negligent Physician's action causes the child's harmed condition—namely, its cognitive disability.

The harm-based approach raises many questions. What, for example, counts as a harmed condition? Is physical unattractiveness or low IQ a harmed condition? Presumably those who argue for this approach wish to avoid the implication that those whose genes make it likely that their offspring would be physically unattractive or have a low IQ would harm their children by causing them to exist (and therefore presumably ought not to have children). There are several options. One would be to equip the notion of a harmed condition with a threshold that would place ordinary ugliness or low intelligence below the threshold but would locate moderately severe

[6] Joel Feinberg, 'Wrongful Life and the Counterfactual Element in Harming,' in his *Freedom and Fulfillment* (Princeton: Princeton University Press, 1992), 6.

mental retardation above it. But in restricting the harm-based approach to cases involving only relatively serious conditions, this revision leaves us without a response in a wide range of cases in which the Non-Identity Problem arises. Consider, for example, a further variant of the Preconception Case in which a couple seek screening for a genetic defect that causes one's child to have an IQ that is roughly 60 points lower than it would otherwise be. As a result of the Physician's negligence, they have a child with an IQ of 90. If the defect had been detected, the man's sperm would have been altered to correct the defect and they would have conceived a different child with an IQ of 150. Surely we want to condemn the Physician's negligence in this variant on substantially the same ground on which we condemn it in the original Preconception Case. But, on the assumption that an IQ of 90 counts as only ordinary low intelligence, the Negligent Physician has not caused the child in this variant to exist in a harmed condition. So, when it incorporates the stipulated threshold, the harm-based approach lacks the resources to explain how the Physician's action has had any bad effect other than the effect on the couple.

Another option is to try to distinguish among the various cases on the basis of differences in causal responsibility. It can be argued, for example, of parents who have a child that is predictably physically unattractive or of low intelligence that, while they are responsible for the child's existence, the fact that the child is unattractive or unintelligent is not attributable to the act that caused the child to exist. There is, in fact, *no* act that causes the child to be unattractive or unintelligent; this is just the way the child is. Thus parents who have an unattractive or unintelligent child do not thereby harm the child. By contrast, consider again the Negligent Physician who administers an untested fertility drug that enables a child to be conceived but also causes the child to develop a serious disease later in life. In this case one and the same act is both a causally necessary condition of the child's existence and the cause of the disease. This act harms the child, though it is not worse for him.

The problem with this response is that it does not seem to divide the cases in the desired way. Reconsider the original Preconception Case. Here the Negligent Physician's action is a causally necessary condition of the retarded child's existence, but it does not seem to be the cause of the retardation, any more than the act of conceiving a predictably ugly child is the cause of the child's ugliness. In each case, that is just the way the child is. So any conception of causal responsibility that allows us to deny that the parents of an ugly child harm the child by causing him to exist seems also to imply that the Negligent Physician in the original Preconception Case does not harm the retarded child. Yet the desire to show that the Physician *does*

harm the child is precisely what motivates people to accept the harm-based approach.

Assume, then, that the harm-based approach accepts that to cause someone to exist with a congenital genetic defect is to harm that person if the defect constitutes or inevitably causes a harmed condition. In that case, it seems that we must revert to the option of distinguishing between congenital conditions that count as harmed conditions and those that are insufficiently serious, or perhaps sufficiently widespread or normal, not to count as harmed conditions. (Otherwise the same objection to causing the retarded child to exist in the Preconception Case will apply in *all* cases of causing people to exist, since *everyone* has congenital characteristics that adversely affect their interests—e.g., I have a constellation of interests having to do with achieving great things in philosophy but am thwarted by deficiencies in native intelligence.) As I noted above, this means that the harm-based approach is at most only a partial solution to the Non-Identity Problem; but even a partial solution may constitute progress.

There remains, however, a further problem. In the Preconception Case, and the other cases with which we are concerned, the child whose existence inevitably involves a harmed condition nevertheless has a life that is worth living. The life is worth living because the goods it contains together outweigh the badness of the harmed condition and its effects within the victim's life. Why cannot we say that the act that harmed the child by causing him to have a harmed condition was not bad because the harmed condition is compensated for by the goods of life that the child would not have had if the act had not been performed? There are many instances in which it is best to harm a person for the sake of the compensating benefits that the harmful act brings to that same person—for example, painful or disfiguring medical procedures that are necessary to save a life. If it is not bad, overall, to cause these harms, why is it bad for the Negligent Physician to cause the harm he causes, which is similarly outweighed? (It might be argued that what makes medical procedures that cause harm permissible is the patient's consent; thus the relevant difference between these cases and the Preconception Case is that in the latter the retarded child cannot consent to accept his retardation as the cost of having the compensating goods of life. But it is easy to imagine cases in which it is best to perform a disfiguring or otherwise harmful operation to save a person's life even when the person cannot consent—e.g., because he is unconscious at the time that the decision to operate must be made.)

The problem here is more serious than it may initially seem. In any case in which a child is caused to exist there is a finite probability that the child will have a congenital defect that will constitute or inevitably cause a harmed

condition. If, in cases of causing a person to exist, a harmed condition cannot be outweighed by the goods that the life will contain—in the sense that it remains worse, other things being equal, to cause the child to exist—then it is difficult to see how the *probability* of a congenital harmed condition can be outweighed by the probability that the life will also contain compensating goods. But, if the probability that a child will have a congenital harmed condition cannot be outweighed by the probability of compensating goods, then it seems that, at least where *expected* effects are concerned and when other things are equal, it is worse to have a child than not to have a child.

One response to this objection is to claim that, in the Preconception Case, the Negligent Physician is responsible for the harmed condition (i.e., cognitive disability) but not for any of the goods that the retarded child's life contains. The goods are attributable to other causes. On this view, there are no benefits attributable to the Negligent Physician's action that are capable of compensating the child for the harm that the action has caused. This response, however, seems untenable. If the Negligent Physician is responsible for the retardation because it is an inherent aspect of the child's nature and is therefore attributable to those causal factors that produced *that* child with *that* nature, he should be equally responsible for those inherent aspects of the child's nature that are good or beneficial for the child. If it makes sense to say that the retardation adversely affects the child's interests, it should also make sense to say that the good inherent aspects of the child's nature positively affect the child's interests.[7] Finally, if to cause the retardation is to cause a harm (or a harmed condition), then to cause the good aspects of the child's nature should be to cause benefits (or beneficial conditions). If all this is right, the Negligent Physician's action not only harms but also benefits the child—not necessarily by causing the child to exist but by causing the child's life to contain certain goods. And, since the child's life is worth living, it is reasonable to suppose that the benefits outweigh the harms.

A second response to the objection is to claim that, in the case of ordinary procreation, the risk of harming the child by causing him to have a congenital harmed condition is outweighed not by any probable compensating

[7] Are a congenitally retarded child's interests adversely affected by the retardation? Consider a nonhuman animal with cognitive and emotional capacities and potentials comparable to those of the retarded child. Are its interests adversely affected by the fact that its cognitive capacities are significantly lower than ours? If not, it is not obvious why one should suppose that the interests of the retarded child are adversely affected by his or her cognitive capacities. One's interests are shaped by one's cognitive capacities. See Jeff McMahan, 'Cognitive Disability, Misfortune, and Justice,' *Philosophy & Public Affairs*, 25 (1996), 3–34.

goods that the child's life might contain but instead by the expected benefits to the parents (and perhaps others in the society) of having the child. It might be thought that this response prevents us from objecting if a couple *deliberately* conceive a child with a congenital harmed condition rather than a normal child. But this worry can be dispelled by noting that parental interests may be sufficiently important to outweigh a *slight* risk of causing a harmed condition without being important enough to outweigh a *high* risk of causing a harmed condition, which there would presumably be if the parents *intended* to cause such a condition. Still, the appeal to parental interests cannot rescue the harm-based approach. For this appeal in effect grants the objection that there is always a presumption against procreation based on the risk of causing a congenital harmed condition—a risk that cannot be offset by the probability of compensating benefits within the life. But it is hard to believe that procreation is, in ordinary conditions, an activity that requires the interests of the prospective parents or of other preexisting persons to tip the balance in favor of permissibility. In ordinary circumstances, there simply is no prima facie objection to or presumption against procreation—or, rather, if there is such a presumption, it derives from current conditions of overpopulation rather from the risk of causing a congenital harmed condition.

The harm-based approach fails because it has no explanation of why an act that is assumed to cause a congenital harmed condition ought not to be done even when it also causes compensating benefits. There is, however, an alternative approach of the same sort—one that identifies a victim—that offers such an explanation. According to this approach, which we may call the *rights-based approach*, there are certain harmed conditions that are sufficiently serious that to cause them constitutes a violation of the victim's rights.[8] Assume that, according to this view, the Negligent Physician in the Preconception Case violates one of the retarded child's rights. (Again this is problematic. It is not obvious exactly what right is supposed to be violated or that the requisite causal connections obtain between the Negligent Physician's action and the relevant aspect of the child's condition.[9] But waive these difficulties.) While the Negligent Physician's action was not worse, or bad on balance, for the child, since the harm it caused is outweighed by compensating goods, that is not a sufficient justification for the action. For,

[8] For a careful exposition of an approach of this sort, see James Woodward, 'The Non-Identity Problem,' *Ethics*, 96 (1986), 804–31.

[9] One way of dealing with this difficulty is to claim not that the Negligent Physician violates the child's rights but that his action causes the child to exist with rights that cannot be fulfilled. Respect for the potential child's rights therefore required that he refrain from doing what would cause the child (and therefore the rights) to exist. Compare Jeff McMahan 'Problems of Population Theory,' *Ethics* 92 (1981): 125.

even if an act is on balance beneficial to a person, or on balance promotes the person's well-being or good, that is in general not a justification for the act if the act also violates the person's rights. Our rights protect us even from certain well-meaning forms of action aimed at our own good. Thus the central objection to the harm-based approach—that it cannot explain why the Negligent Physician's action is wrong if the harm it causes is outweighed by compensating benefits—is met by the rights-based approach.

But the rights-based approach faces other objections. Imagine a disability—*condition X*—that is not so bad as to make life not worth living but is sufficiently serious that to cause someone to exist with condition X would be, according to the rights-based approach, to violate that person's rights. One objection to the rights-based approach is that, if it is wrong to cause someone to exist with condition X, it should also be wrong to save someone's life if the only way of doing so would also cause the person to have condition X and it is not possible to obtain the person's consent to being saved in this way. Suppose, for example, that a late-term fetus (which we may assume would be numerically identical with the person into whom it would develop) contracts a disease that requires a certain treatment in order to survive but that the treatment inevitably causes condition X. Whether or not there is a strong moral reason to save the fetus for its own sake, it seems intuitively clear that it would not be *wrong* to treat the fetus, thereby saving its life. But saving it involves causing it to have condition X and thus, apparently, violates its rights. If it is not permissible to violate a right on the ground that the act that violates the right on balance benefits the right-bearer, then it seems that the rights-based approach implies that it *would* be wrong to save the fetus.

The proponent of the rights-based view may reply that this is a case involving a conflict of rights. While the fetus has a right not to be caused to have condition X, it also has a right to be saved.[10] And in this case the right to be saved, being more important, overrides the right not to be caused to have condition X. This reply assumes, however, that priority between the two rights is determined by the comparative strengths of the interests they protect. But the strength of a right does not vary proportionately with the strength of the interest it protects (assuming that it protects an interest at all). The importance of any interest it might protect is only one of a number of factors that contribute to determining the strength of a right. Among the more important determinants is whether the right is positive or negative. The right to life, or the right not to be killed, and the right to be saved both protect the same interest: namely, the interest in continuing to life or in

[10] I owe this response to Frances Kamm.

avoiding death. But the right not to be killed is a negative right and is thus held, by theorists of rights, to be considerably stronger, other things being equal, than the right to be saved. But if negative rights are in general considerably stronger than corresponding positive rights, it is at least arguable that the negative right not to be caused to be disabled is stronger, or more stringent, than the positive right to be saved.

Let us suppose, however, if only for the sake of argument, that it is true that the late-term fetus's right to be saved overrides its right not to be caused to have condition X, so that the rights-based approach does not imply that it would be wrong to save the fetus. Now consider a parallel case involving an *early-term* fetus. The fetus has a disease that will rapidly be fatal unless it is treated; but the treatment causes condition X. Assume that the claim that I noted earlier is correct—namely, that individuals such as you and I do not begin to exist until our organisms acquire the capacity to support consciousness and mental activity. If that is right, an early-term fetus does not support the existence of an individual of the sort that you and I essentially are. There is no one there to have a right to life or a right to be saved. Hence, according to the rights-based approach, there is a strong reason *not* to treat the fetus, since treating it would violate the right of the later person not to be caused to have condition X, but no countervailing rights-based reason to treat it. The rights-based approach therefore implies that it would be *wrong*, other things being equal, to treat the fetus.

This seems an implausible result. But what is even more implausible is that the rights-based approach distinguishes morally between the case of the late-term fetus and the case of the early-term fetus, claiming that one may treat the former but not the latter. As I indicated, the approach *may* imply that it is wrong to treat the late-term fetus as well. If so, the approach would avoid the embarrassment of treating the two cases asymmetrically. But the claim that it is wrong to treat the late-term fetus is itself quite implausible. There is, of course, a way around the dilemma, which is to reject the view that we do not begin to exist until the fetal organism develops the capacity to support consciousness. If we begin to exist when the human organism begins to exist, shortly after conception, it may be defensible to claim that in both cases the fetus has a right to be saved and that this makes it permissible to treat the fetus despite the fact that doing so infringes its right not to be caused to be disabled. But the supposition that even early fetuses have a right to be saved from death is quite a radical view, with implications for abortion and other issues that many will be reluctant to accept.

If I am right about the metaphysics, the case of the early-term fetus involves a choice between causing a child to exist with a disability and allowing it to be the case that the child fails to come into existence. It is *not* a

feature of the case that, if the early-term fetus is untreated, a different, normal child (i.e., without either the disease or condition X) will be caused to exist instead. Cases of this sort are helpful in testing the plausibility of the harm-based and rights-based approaches. For these approaches hold that what is fundamentally objectionable about causing a person to exist with a congenital disability (i.e., one that constitutes a harmed condition or necessarily causes a right to remain unfulfilled) is found in the inherent condition of the person, not in anything extrinsic to the person's life. Thus their plausibility can best be tested by reference to cases in which the only conceivable objectionable features are intrinsic to the life of a person caused to exist with a disability. In other cases, such as the Preconception Case, in which a person is caused to exist with a disability *and* there was the alternative of causing a normal child to exist instead, these approaches may yield the intuitively correct judgment but for the wrong reason. For it may be that the comparative dimension to the case—i.e., that the Negligent Physician causes a disabled child to exist *rather than* a normal child—is an essential part of the explanation of why it is objectionable to cause the disabled child to exist.

To test the approaches that identify a victim, we should therefore consider cases that lack this comparative dimension. Imagine, then, a situation in which *any* child one might cause to exist would have a congenital harmed condition, or a right that would necessarily remain unfulfilled. It is not possible, in the circumstances, to cause a normal child to exist instead. Would it be wrong for a couple who wish to have a child to conceive a child in these circumstances? Most people believe that, provided that the child's life would be worth living and that the motives of those who would cause the child to exist would not be discreditable, it would not be worse, or bad, or wrong (other things being equal) to cause the child to exist. This is not just an intuition. The reason that it is not bad to cause the child to exist is, as I suggested in discussing the harm-based approach, that the goods that the child's life contains compensate for the presence of the harmed condition, without which the child would not exist. Thus the fact that the harm-based and rights-based approaches imply that it *would* be wrong to cause the child to exist constitutes a serious objection to them.[11]

What we need is an account that explains why it is objectionable to cause a disabled child to exist when it would be possible to cause a normal child to exist instead (as in the Preconception Case) but accepts that it is not bad, and is thus permissible if other things are equal, to cause a child to exist with

[11] For an intricate and detailed critique of the harm-based and rights-based approaches, see Parfit, 'Comments,' 854–62.

the same disability when any child one might cause to exist would necessarily have that disability. It is difficult to find an approach of the victim-based type that does both these things since these approaches do not locate the objection to causing a disabled child to exist in factors that are comparative or in any way extrinsic to the condition of the child. There is, however, one approach of this sort that has a certain amount of promise. This account invokes the notion of a *restricted life*—a notion introduced by Kavka in his influential and important paper on the Non-Identity Problem.[12] Kavka defines a restricted life as 'one that is significantly deficient in one or more of the major respects that generally make human lives valuable and worth living.' He goes on to note, however, that 'restricted lives typically will be worth living, on the whole, for those who live them.'[13] I will use Kavka's suggestive term in a slightly different way to refer to a life that is objectively not worth living but is subjectively tolerable, and may indeed be overall enjoyable to the individual whose life it is. Such a life is, I will say, subjectively worth living but objectively not worth living. (I put aside the question whether there could be a life that was objectively worth living but subjectively not worth living.) As an example, consider the life of Adolf Hitler. There is reason to believe that Hitler was, during most of his adulthood, abundantly satisfied with his life. Judged by the usual standards, he was a reasonably happy man. His life was therefore subjectively worth living: *he* found it well worth living. But was his life *objectively* worth living? Was this in reality a good life for him to have—better, at least, than no life at all? It is plausible, I think, to claim that Hitler's adult life was a *dreadful* life—not just in its effects on others but dreadful *for him* (even though he himself failed to recognize this). This is not the kind of life that it could be good for anyone to have. It would have been better *for Hitler* if he had died in his twenties.

How does the notion of a restricted life help with the Non-Identity Problem? Assume that the retarded child in the Preconception Case has a restricted life. This explains why it was bad that the Negligent Physician's action resulted in the child's existence: the child's life is not worth living; it is objectively bad *for the child* to exist with that sort of life. If the child's life is genuinely restricted, the goods that it contains do not, on balance, compensate for the child's harmed condition. This also supports the claim that the Negligent Physician owes the child compensation. For the Physician's negligence was culpable and had a victim: the child, for whom the Physician's action was bad.

[12] Gregory S. Kavka, 'The Paradox of Future Individuals,' *Philosophy & Public Affairs*, 11 (1982), 93–112.

[13] Ibid. 105.

This explanation is, however, essentially noncomparative: it does not mention the alternative possible outcome in which a normal child would have existed. It focuses entirely on the intrinsic features of the retarded child's life. How, then, can it explain the permissibility of causing a child with a life like this to exist when it would not be possible to cause a normal child to exist instead? This approach must, it seems, claim that there is a serious prima facie reason not to have such a child—namely, that it would be objectively bad for the child. Yet, although a life that is objectively not worth living is bad, it is not nearly so bad if it is subjectively worth living as it would be if it were also subjectively not worth living. For a life that is objectively not worth living but is nevertheless subjectively worth living is not experienced as a burden by the person whose life it is. Thus the moral presumption against causing a person to exist with a restricted life may be overridden by countervailing considerations that are considerably weaker than those that would be required to override the much stronger presumption against causing a person to exist with a life that would be both subjectively and objectively not worth living. Assuming, then, that the desire to have a child has a certain normative force (e.g., that it is supported by a right of procreation), it might be that the desire of a couple to have a child could be sufficient to outweigh the harm they would do to the child by causing it to exist with a restricted life. But this same desire would be insufficient to justify causing a child to exist with a restricted life when it would be possible to have a normal child instead. For the reasonable desire to have a child could be satisfied by having the normal child. There would have to be some *other* reason to justify doing what would cause a child with a restricted life to exist rather than a normal child. And in the ordinary circumstances of life it is doubtful that there could be a reason sufficiently strong to justify the harm to a child with a restricted life.

Even when it would be permissible for a couple to cause a child to exist with a restricted life, the child would have a claim to compensation comparable in force to that which the retarded child has against the Negligent Physician in the Preconception Case. In practice this means that a couple that chose to have a child with a restricted life would be morally required to make sacrifices for the child that would not be part of the normal burden of child-rearing. This, however, seems entirely plausible. Whether such parents would owe as much as the Negligent Physician depends on whether the fact that he is at fault compounds his liability.

The appeal to the notion of a restricted life thus has a certain promise. But it nevertheless faces serious objections. It may be objected, for example, that the killing of people judged to have restricted lives could be justified as euthanasia. This, however, is not a serious concern. For the morality of killing is not governed solely by considerations of harm and benefit. Even though there is a sense in which it would be better for a person with a restricted life to die rather than continue to live, it certainly does not follow that it would be permissible to kill that person against his or her will. It has to be conceded, however, that the notion of a restricted life is an exceedingly dangerous one; for it asserts the possibility that others could know that one's life was not worth living even if one were oneself convinced that it was worth living. As a matter of principle this in fact seems to be right: it is possible to believe that one's life is worth living when in fact it is not. But surely this occurs very rarely and in most cases one's judgment about whether one's own life is worth living is, if not authoritative, then at least so nearly infallible that it would be the height of presumption for another person to dispute it.

This observation reveals the central weakness of the appeal to the notion of a restricted life—that there are scarcely any plausible instances of lives of this sort.[14] Perhaps the most plausible examples are of people whose lives, while enjoyable, are utterly morally debased. But these cases are largely irrelevant to the morality of causing people to exist because of the impossibility, at least at present, of predicting before a life begins that it will be morally degraded. Among predictable conditions, it is difficult to identify *any* that clearly make a life objectively not worth living without making it subjectively not worth living as well. Perhaps the most plausible candidate is severe congenital cognitive incapacity. Loren Lomasky contends that, 'were one condemned . . . to remain a child throughout one's existence, or to grow in bulk without simultaneously growing in the capacity to conceptualize ends and to act for their sake, it would be a personal misfortune of the utmost gravity.'[15] The idea that the severely retarded are appropriately viewed as permanently infantile suggests that a life in this condition may be objectively degraded or unworthy, even if it is subjectively tolerable. It is, however, difficult to reconcile this judgment with the commonly accepted assumption that the lives of nonhuman animals with comparable cognitive capacities may be worth living and are certainly not objectively degraded simply because they are not guided by the exercise of our higher cognitive capacities.[16]

The Non-Identity Problem arises in a large number of cases. Since it is

[14] The limited applicability of Kavka's notion of a restricted life to the Non-Identity Problem is noted by Derek Parfit in his reply to Kavka, 'Future Generations: Further Problems,' *Philosophy & Public Affairs*, 11 (1982), 120–1.

[15] Loren Lomasky, *Persons, Rights, and the Moral Community* (New York: Oxford University Press, 1987), 202.

[16] Cf. 'Cognitive Disability, Misfortune, and Justice.'

difficult to think of a single case in which a predictable condition causes a person to have a restricted life, it is safe to conclude that the appeal to the notion of a restricted life cannot solve this problem.[17]

III THE IMPERSONAL COMPARATIVE APPROACH

Other writers have discussed cases with the same structure as the Preconception Case under the heading of 'wrongful life.' Some have contended that the fundamental objection in these cases to causing a child to exist with a disability is that this gratuitously increases the amount of evil, or that which is bad, that the world contains. Joel Feinberg, for example, claims that the agent in a case such as the Preconception Case 'must be blamed for wantonly introducing a certain evil into the world, not for harming, or for violating the rights of, a person.' He then goes on to elucidate the nature of the evil when he observes that one could make the willful creation of a disabled child a criminal act on the ground that 'the prevention of unnecessary suffering is a legitimate reason for a criminal prohibition.'[18] These remarks are echoed by John Harris, who writes: 'What then is the wrong of wrongful life? It can be wrong to create an individual in a harmed condition even where the individual is benefited thereby. The wrong will be the wrong of bringing avoidable suffering into the world, of choosing deliberately to increase unnecessarily the amount of harm or suffering in the world.'[19]

To say that some instance of suffering is 'unnecessary' or 'avoidable' is to imply that it is wrong to cause that suffering. But what exactly does it mean to say that suffering is unnecessary or avoidable? In one sense, it means simply that the suffering could have been avoided. But that is true of *all* human suffering, since it has always been possible for people simply to stop procreating. Thus those unguarded forms of Negative Utilitarianism that call simply for the minimization of suffering have notoriously been accused of implying that it is wrong ever to cause a sentient being to exist. But this is surely not the sense of 'unnecessary' intended by Feinberg and Harris. The normal implication of the claim that some instance of suffering is unnecessary is that the suffering is not instrumental to or a necessary accompaniment of some greater good for the person who experiences the suffering. Thus suffering that is *not* unnecessary is suffering that has to occur if certain compensating goods are to be had by the sufferer. In this sense, the foreseeable suffering that any life that is worth living will inevitably contain is not unnecessary. But then this applies equally to the foreseeable suffering within the lives of the congenitally disabled, provided that their lives would be worth living. Their suffering is not unnecessary in this second sense.

[17] It may be easier to find examples in which a *part* of a life is restricted. Imagine that a person whose life has hitherto been devoted to intellectual pursuits suffers brain damage and becomes a contented idiot. Her subsequent life may be subjectively tolerable from her present point of view but objectively not worth living in the light of values that she autonomously embraced prior to the loss of her cognitive competence.

[18] Feinberg, 'Wrongful Life and the Counterfactual Element in Harming,' 27 and 28.

[19] John Harris, *Wonderwoman and Superman: The Ethics of Human Biotechnology* (Oxford: Oxford University Press, 1992), 90.

Feinberg and Harris must therefore be invoking a third sense in which suffering may be unnecessary. Consider again the Preconception Case. Whatever suffering the child experiences as a result of the retardation is not unnecessary for the compensating goods of *that* life. But it is unnecessary for goods of the same type—indeed a greater quantity of those goods—within a *different* life that might have been caused to occur instead. The objection urged by Feinberg and Harris therefore takes an impersonal, comparative form. For it is not concerned with effects for better or worse on any particular individual but with the comparison between the possible effects on one possible individual with those on another. The objection to causing the retarded child to exist is that it was possible to cause a *different* child to exist whose life would have contained at least as much good but less of what is bad—in particular, less overall suffering. It is in this impersonal sense that the retarded child's suffering is unnecessary.

A more precise articulation of this sort of approach has been formulated by Parfit in the following principle: 'If in either of two possible outcomes the same number of people would ever live, it would be worse if those who live are worse off, or have a lower quality of life, than those who would have lived.'[20] Call this the *Impersonal Comparative Principle*. Notice that it is explicitly restricted to what Parfit calls 'Same-Number Choices'—that is, cases in which the same number of people would exist in all the possible outcomes of a choice between acts. Note furthermore that, being impersonal, the Impersonal Comparative Principle is consistent with the No-Difference View, which asserts, in effect, that the correct principle of beneficence must take a fully impersonal form. Finally, notice that this principle presupposes that possible people count morally and must be taken into account in moral deliberation.

The Impersonal Comparative Principle has a distinct advantage. It does seem, intuitively, that the morality of causing a disabled child to exist is affected by whether or not it would be possible to cause a normal child to exist instead. The Impersonal Comparative Principle captures this. Thus it

[20] *Reasons and Persons*, 360. Kavka proposed a similar principle but rejected it in 'The Paradox of Future Individuals,' 99–100.

condemns the Negligent Physician's action in the Preconception Case because the normal child who might have existed would have been better off than the retarded child is. But it does not condemn a couple for having a child with the same disability provided that the child's life is worth living and that any child they might have would also have that disability. It does not condemn such a couple because it has *no* implications for their choice. Their choice is between having a disabled child and having no child. It is therefore not a Same-Number Choice but what Parfit calls a 'Different-Number Choice'—that is, a case in which different numbers of people would exist in some of the possible outcomes of a choice between acts.

Does the Impersonal Comparative Principle support the intuition that the Negligent Physician in the Preconception Case owes compensation to the retarded child? Here is an argument for the claim that it does. According to the Impersonal Comparative Principle, the Physician has a moral reason *ex ante* to ensure that a normal child exists rather than a disabled child. Indeed, his general reason to bring about the better outcome is strengthened in this case by his professional commitment. Presumably he would even have been required, if necessary, to accept certain costs in order to ensure the conception of a normal child rather than a retarded child. (The extent of the cost he should accept in order to ensure the better outcome is of course limited. If, for example, the cost to him of ensuring the conception of a normal child rather than a disabled child would be as great as the cost to a couple of being unable to have a child, then it might be permissible for him to allow the conception of a disabled child. But it is hard to imagine circumstances in which personal costs this great would be required from a physician in order to ensure the conception of a normal rather than a disabled child.) Let us stipulate that the Negligent Physician in the Preconception Case would have been required to accept costs up to amount x in order to ensure that the couple would conceive a normal rather than a disabled child. If that is true, it seems reasonable to suppose that, since his negligence has brought about the worse outcome, he should be required *ex post* to pay costs at least up to amount x in order to repair the result of his fault. In particular, he should be required to pay up to amount x, if necessary, to try to make the disabled child's life as good as the normal child's life would have been. If the compensation could succeed in benefiting the disabled child to that extent, this would cancel the bad effect of his previous action.

This argument is vulnerable to several objections. First, it is not plausible to suppose that there are grounds for liability whenever the Impersonal Comparative Principle implies that it was worse to cause some person to exist. For the Impersonal Comparative Principle implies that it is worse, other things being equal, to cause a person to exist whenever it would be possible to cause a different, better-off person to exist instead. Thus it implies that it would be worse to cause a normal person to exist if it would be possible to cause a person with an unusually high capacity for well-being to exist instead. But if, in these circumstances, one were to cause the normal person to exist, it is implausible to suppose that this would make one liable to compensate that person for being worse off than some extraordinary possible person might otherwise have been.

Second, the case for compensation depends on the availability of a better alternative. The Impersonal Comparative Principle does not imply that there is a reason not to cause a disabled child to exist when there is no possibility of causing a better-off child to exist instead. Hence there is no reason in these circumstances for an agent to accept costs *ex ante* to avoid causing a disabled child to exist and no basis for a claim to compensation *ex post*. But now imagine two equally disabled children, only one of whom was caused to exist in conditions in which a normal child could have been caused to exist instead. Although both have the same disability, only this child can claim compensation. But it may seem unfair to deny the other compensation just because there was no possibility of causing a normal child to exist in his place.

Finally, and most importantly, recall that the reason that the Impersonal Comparative Principle holds that the Negligent Physician's action was worse is not that it harmed or wronged the retarded child. His offense was instead *impersonal*. But, if the original action was objectionable for impersonal reasons, the reason to redress the situation should be impersonal as well. There is, in other words, no reason why the remedy—that is, the action aimed at canceling the bad effect—should benefit the disabled child. After all, that child is not, according to the Impersonal Comparative Principle, a victim of the Negligent Physician's action.

Despite initial appearances, therefore, the Impersonal Comparative Principle provides no basis for liability on the part of the Negligent Physician to compensate the disabled child. This may or may not constitute an objection to the principle. For it is unclear whether, in the Preconception Case, the child in fact deserves compensation. The child may deserve special compensation through relevant mechanisms of social redistribution simply for being badly off—either in absolute terms or relative to the norms of the society. In this respect the child is on a par with others who are badly off through no fault of their own. There is no reason why the Negligent Physician in particular should be required to do more than anyone else to help the child.

But, while it is not clear whether the disabled child in the Preconception Case deserves compensation, it *is* clear that the disabled child in the Prenatal Case deserves compensation and that it is the Negligent Physician who is

morally (and legally) liable to pay it. Recall, however, that according to the No-Difference View, the objection in the Prenatal Case to the Negligent Physician's causing the child to be disabled rather than normal is the *same* as the objection in the Preconception Case to the Negligent Physician's causing a disabled child to exist rather than a normal child. The objection in the Preconception Case is impersonal in character; therefore the objection in the Prenatal Case must also be impersonal—which, of course, is exactly what the Impersonal Comparative Principle implies, since it treats the two cases in exactly the same way. Indeed, according to the generalized No-Difference View, the whole of the morality of beneficence is to be explained in impersonal terms. Person-affecting principles may often yield the right answers but they *never* provide the correct explanation, which is always impersonal. If it is worse to perform some act, it is not because the act is bad or worse *for somebody*; there are never any victims in the relevant sense. Notice, however, what this implies. If the objection to the Negligent Physician's action in the Prenatal Case is impersonal, then there can be no more basis for liability here than there is in the Preconception Case. Indeed, if the generalized No-Difference View is correct, then there can never be any basis for liability to compensate an individual for harm that one has done to that individual. Or at least this is true within the area of morality concerned with beneficence, or well-being. Parfit leaves it open that there may be areas of morality governed by respect for rights, or other considerations beyond the scope of beneficence. But if, as Parfit assumes, such cases as the Prenatal Case, in which one person's negligence causes another to suffer a serious disability, come within the morality of beneficence, then it seems that the areas governed by rights cannot be more than tiny provinces at the periphery.

Most of us firmly believe that, in the Prenatal Case, the Negligent Physician owes compensation to the child he has caused to be disabled rather than normal. That the Impersonal Comparative Principle seems incapable of supporting this belief is a serious objection to it, on the assumption that the No-Difference View is true. If the Prenatal Case were outside the proper scope of the Impersonal Comparative Principle, there would be no problem. But the No-Difference View holds that there is no relevant difference between the Prenatal Case and the Preconception Case, that the objection to the Negligent Physician's conduct is therefore the same in each, and that that objection is provided by the Impersonal Comparative Principle.

That the Impersonal Comparative Principle cannot account for the Negligent Physician's liability in the Prenatal Case is only one of many problems it faces. Here is another. As it is stated, the principle refers only to people. But there is no obvious reason why it should not apply to nonpersons as

well. But, when extended in this way, it implies that it is worse, and therefore presumptively objectionable, to breed one's dog rather than to have a child, if one cannot do both. For to breed the dog would be to cause a worse-off rather than a better-off individual to exist.[21] And it would be worse to breed one's lizard than to breed one's dog, if one could not do both. And so on.

These implications are implausible. There are several ways that a defender of the Impersonal Comparative Principle might seek to avoid them. One would be to appeal to side-effects—for example, by arguing that, because of overpopulation, causing a person to exist has such bad side-effects that on balance it would not be worse to breed one's dog instead. But this response is inadequate. It is not because of human overpopulation that it is permissible to breed one's dog. And in any case the principle still implies that it *was* worse, before overpopulation arose, to breed one's dog rather than to have a child.

A second possible response is to note that the Impersonal Comparative Principle, as stated by Parfit, refers only to what is *worse*, not to what is *wrong*—that is, it concerns only the evaluation of outcomes, not what one ought or ought not to do. Therefore, it is only if the principle is conjoined with something like Act Consequentialism that it has implausible implications about procreation and breeding. Again, however, this response is inadequate. For the Impersonal Comparative Principle must be conjoined with *some* principle that explains how considerations of consequences should guide our action; otherwise its utility will be extremely limited when applied to cases like the Preconception Case. It is fairly obvious in that case that the Negligent Physician brings about the impersonally worse of two possible outcomes. What is important is the further claim that this is what explains why his action was morally objectionable or wrong, other things being equal. And it is reasonable to expect that any action-guiding principle that, when conjoined with the Impersonal Comparative Principle, implies that it is wrong to bring about the worse of the two outcomes in the Preconception Case will also imply that it is wrong to bring about the worse of the two outcomes when the choice is between having a child and breeding one's dog.

Perhaps the most plausible response to this challenge is to restrict the scope of the Impersonal Comparative Principle so that it applies only to cases involving lives of the same kind. Suitably restricted, it would imply that it is worse to cause the worse-off of two possible people to exist, and worse to cause the worse-off of two possible dogs to exist; but it would have

[21] Here I follow Robert Merrihew Adams, 'Must God Create the Best?', *Philosophical Review*, 81 (1972), 329.

nothing to say about whether it is worse to cause a dog to exist rather than a person. While I am skeptical that a principled rationale for such a restriction could be found, I cannot exclude the possibility.

The deepest problems for the impersonal conception of beneficence that is required by the No-Difference View emerge when we try to extrapolate beyond the Impersonal Comparative Principle to a principle that covers not only Same-Number Choices but also Different-Number Choices. Among those choices in which different people exist in the different possible outcomes, Different-Number Choices are significantly more common than Same-Number Choices. It is therefore essential to have a principle that covers those choices. There is, however, a formidable obstacle to extending the impersonal approach so that it applies in these cases. The extended principle must surely imply, as the Impersonal Comparative Principle does, that an outcome is worse if the people who exist in that outcome are worse off than the people who would exist in an alternative outcome. But when different numbers of people exist in the different outcomes, it becomes very difficult to determine which group is better off than the others. One has to weigh the number of lives, or perhaps the overall quantity of life, against the overall quality of life. And one has to determine how to measure the overall quality of life in a group in which individual lives may vary considerably in overall quality. Is the group with the best overall quality of life the one with the highest average quality of life, the highest maximum, or perhaps the lowest minimum? Should the measurement of overall quality of life take into account the relative levels of equality in the quality of life within the different groups? And if so, how is equality itself to be measured?

These and other problems are explored with tremendous subtlety and ingenuity in Parfit's book, *Reasons and Persons*. He assigns the label 'Theory X' to the theory that would plausibly extend the Impersonal Comparative Principle so that it would cover Different-Number Choices. While he states a variety of requirements that Theory X would have to satisfy in order to be acceptable, he confesses his own inability to discover the content of the theory. He concludes, however, with an expression of optimism: 'Though I failed to discover X, I believe that, if they tried, others could succeed.'[22]

I believe that there is reason to doubt this. Theory X must take an impersonal form: it must presuppose that the fact that an act is bad or worse *for someone* cannot be part of the fundamental explanation of why its effects are bad or why the act itself is wrong. Because of this, I suspect that any candidate for Theory X will have implications that undermine its credibility. In order to try to substantiate this suspicion, I will indicate what I think

[22] *Reasons and Persons*, 443.

some of these implications are. I must acknowledge, however, that I cannot demonstrate that Theory X will have these implications. As yet there is no Theory X; therefore neither I nor anyone else can say what its implications might be. My claim can only be that it is difficult to see how any candidate for Theory X can avoid the implications to which I will call attention.

Let us revert to a problem mentioned earlier in Section I: the problem of abortion. On the assumption that a new individual of our kind does not begin to exist until some time during the second half of pregnancy, the choice between having and not having an early-term abortion is a Different-Number Choice: the number of people who will ever exist if one has the abortion will be different from the number who will exist if one does not. This is, however, a very simple Different-Number Choice. Consider:

The Early-Term Abortion
A woman is in the very early stages of pregnancy. If she continues the pregnancy, the child she has will have a life that is well worth living. It would be better for her and her partner, however, if she has an abortion. But, because the society in which they live is underpopulated, the abortion would also have certain bad effects on other people. Assume that these various good and bad effects on preexisting people counterbalance one another—that is, they cancel each other out. The couple decide to have the abortion. Overall this is not worse for the people who ever exist.

Suppose we want to know which of the two possible outcomes is better *impersonally*. The complications mentioned earlier that typically make it so difficult to determine which of two different-sized groups is better off simply do not arise in this case. In the actual outcome, a certain number of people exist. If the abortion had not been performed, exactly those same people would have existed and overall their collective level of well-being would have been the same. The only difference is that in the second outcome there would have been one additional person whose life would, we may assume, have been worth living. (There are instances in which, when one thing that is good when taken by itself is added to a second thing that is also good by itself, the result is a decrease in the degree of goodness of the second thing. Nothing like this would occur if, in the Early-Term Abortion, the abortion were not performed and a new person were added to the existing population.) But if, from an impersonal point of view, the two outcomes differ only in that one contains an additional life in which the good elements outweigh the bad, then it seems that the outcome with the additional good should be better impersonally.

One might arrive at the same conclusion by a slightly more circuitous route. Let us define three outcomes: having a Happy Child, having a Less Happy Child, and having No Child. According to the Impersonal Comparative Principle, having a Less Happy Child is worse than having a Happy

Child, other things being equal. This is not because having the Less Happy Child would be bad in itself; it is just that having a Happy Child contributes more to making the world better. But if having a Happy Child is better than having a Less Happy Child because it adds more good to the world, then it seems that having a Happy Child must also be better than having No Child, other things being equal, and for the same impersonal reason.

In the Preconception Case, the Negligent Physician causes the couple to have a Less Happy Child rather than a Happy Child. In the Early-Term Abortion, the couple have No Child rather than a Happy Child. From an impersonal of view, the latter should be as objectionable as the former. If we conclude that the Negligent Physician *ought not* to have caused the less good outcome rather than the better one, and for reasons that are impersonal, it may be difficult to avoid the conclusion that the woman in the Early-Term Abortion *ought not* to have had the abortion. It seems that Theory X, which will extend the claim of the Impersonal Comparative Principle so that it covers Different-Number Choices, may imply that abortion is wrong. The impersonal approach to the Non-Identity Problem thus not only threatens a powerful argument in favor of the permissibility of abortion (as I suggested in Section I) but also supports an argument against the permissibility of abortion.

Indeed the problem runs deeper than this. The objection to abortion that seems to be implied by the impersonal approach cannot, of course, be that abortion is murder, that it harms the fetus, or that it is against the fetus's interests. It is simply that abortion prevents the existence of a person whose existence would make the outcome better in impersonal terms. But this is equally true of the use of contraception and indeed of any choice that results in abstention from procreation. To the extent that abortion is objectionable from an impersonal point of view, these other forms of behavior must be objectionable as well, other things being equal, and for the same reason.

The common sense view is of course entirely different. Most of us believe that there is no moral reason to cause a person to exist just because the person's life would be worth living—that there is no reason, other things being equal, to have a Happy Child rather than No Child. But we also believe that, *if* one is going to have a child, one has reason, other things being equal, to have a Happy Child rather than a Less Happy Child. The moral reason for having a Happy Child is conditional on a prior determination to have a child. Thus we believe that it is permissible to have No Child rather than a Happy Child even though it is wrong to have a Less Happy Child rather than a Happy Child, other things being equal. I have suggested, however, that it is difficult to see how this set of beliefs could be defensible within an impersonal conception of beneficence.

Parfit has suggested an analogy that might be thought to show how these beliefs could be consistent.[23] Suppose, Parfit writes, that I have three alternatives:

A: at some great cost to myself, saving a stranger's right arm;
B: doing nothing;
C: at the same cost to myself, saving both the arms of this stranger.[24]

Most of us believe that, if these are the alternatives, it is permissible to do B—that is, to save neither arm. But, if one has decided to help the stranger, it would be wrong to do A—that is, to save one arm rather than two. If one has decided to accept a certain cost to help the stranger, and the cost will be the same whether one saves one arm or both, it would be perverse not to do what would achieve the greater good. In short, while there is no duty to do C rather than B, there is a duty to do C rather than A.

Now alter the values of the variables so that one's alternatives are:

A: having a Less Happy Child;
B: having No Child;
C: having a Happy Child.

Again the common view is that, if these are the alternatives, it is permissible to do B—that is, it is permissible not to have a child. But, if one has decided to have a child, it would be wrong to do A—that is, to have a Less Happy Child rather than a Happy Child. As in the first set of alternatives, there is no duty to do C rather than B, though there is a duty to do C rather than A. This is the common sense conception of the morality of procreation. The parallel with the first set of alternatives suggests that this conception is defensible and hence that I was mistaken to claim that, if C is better than A, it must also be better than B.

This counterargument fails, however, for the two sets of alternatives are not in fact parallel. Once parallelism is established, the comparison between them supports rather than refutes my claim. In the first set of alternatives, C is the best outcome, impersonally considered. There is also a strong moral reason to do C rather than B. It is only because there is a great cost to the agent attached to C that it is permissible to do B rather than C. If we

[23] Parfit is discussing his own revised version of a principle suggested but ultimately rejected by Kavka. Thus he does not himself employ his analogy the way that I do here and my critique of the analogy is not directed against his discussion. See his 'Future Generations: Further Problems,' 127–32.

[24] Ibid. 131.

subtract the stipulations about costs from the first set of alternatives, so that this set becomes analogous to the second, one would be required to do C rather than B (which is the conclusion that I assume is implied by the impersonal approach in the second set of alternatives). Alternatively, if one adds parallel stipulations about cost to the second set of alternatives, common sense intuitions may be upheld but the explanation of why it is permissible to have No Child is no longer that this outcome is not worse than having a Happy Child. The explanation instead appeals to considerations of cost, which are extraneous to the impersonal evaluation of the outcomes.[25] Indeed, to maintain parallelism with the first set of alternatives, it must be granted that there is a strong moral reason to have a Happy Child rather than No Child. This reason is overridden only by considerations of cost to the agent.

So the point still stands: it seems that Theory X will imply that it is better for a Happy Child to exist than for No Child to exist and consequently that there is a moral reason to have a Happy Child rather than No Child. While common sense resists this claim, it is not obviously wrong. But there is worse to come. For, from an impersonal point of view, there seems to be no fundamental difference between starting a life and extending a life. Provided that each would be worth living, one's reason to create a new life is the same as one's reason to extend an existing life: namely, that doing either makes the outcome better by causing there to be more of what is good, or that which makes life worth living. This suggests that, other things being equal, Theory X will imply that there is as much reason to cause a new person to exist as there is to save a person's life. Indeed, since the outcome of saving a life contains only a *part* of that life, whereas the outcome of causing a person to exist contains the *whole* of a life, it will normally be *better*, other things being equal and from an impersonal point of view, to cause a person to exist than to save a person's life.

This is very hard to believe. But there is more. Accounts of the morality of beneficence that are impersonal in character tend to treat as irrelevant certain aspects of an agent's mode of agency. They tend, for example, to deny that there is any moral significance to the distinction between doing and allowing, or to the distinction between effects that are intended and those that are foreseen but unintended. While there is no necessary incompatibility between an impersonal theory of beneficence and claims about the significance of agency that are essentially deontological in character, it is nevertheless natural that a theory that evaluates outcomes impersonally should also take an impersonal view of agency. If the identity of the beneficiary or victim of an act makes no difference to the morality of the act, it should not be surprising if neither the identity of the agent nor his or her mode of agency matters either. Thus many writers who accept an impersonal theory of beneficence deny that there is any fundamental or intrinsic difference between failing to save a person and killing a person. But, if this is right, and if there is also no fundamental difference between saving a person and causing a person to exist (or between not saving a person and not causing a person to exist), it follows that there is no difference, other things being equal, between killing a person and failing to cause a person to exist. From an impersonal point of view, both are bad for the same reason: the outcome is worse because it contains less good—less good than it would have contained had a person with a life worth living continued to exist or been caused to exist. Indeed, failing to cause a person to exist will, other things being equal, be *worse*, for the same reason that it is normally impersonally worse than failing to save a person.

[25] Parfit's own discussion of the parallels between the first and second sets of alternatives explicitly appeals to considerations of cost to the agent in order to explain why 'most of us . . . have no duty to have unwanted children.' Ibid. 128.

It might be argued that, even if Theory X has these implications, this does not show that it is unacceptable. For Theory X is an account of beneficence only, and there is more to the morality of killing than considerations of beneficence. Killing may be specially objectionable, for example, because it involves a violation of rights. But, if this defense works at all, it applies only to the comparison between killing and failing to cause a person to exist. For it is implausible to suppose that the morality of saving lives lies outside the scope of beneficence.

A second response might be to argue that it is compatible with a wholly impersonal conception of beneficence to suppose that there is a moral asymmetry between harms and benefits, or between suffering losses and forgoing gains, or something of the sort. If there is such an asymmetry, then even within the morality of beneficence killing a person is worse than failing to cause a person to exist, since killing involves harm or loss while the failure to cause a person to exist involves only the absence of benefit or gain. This, however, is a mistake. The harm of death consists primarily if not exclusively in the loss of the benefits of continued life. Death and the failure of a person to come into existence involve the same sorts of loss from the impersonal point of view.

Finally, even if there are dimensions to the morality of killing beyond the evaluation of outcomes (and I believe that there are), Theory X seems to get even the evaluation of outcomes wrong. If we compare an act of killing with a failure to cause a person to exist, it seems that the *outcome* of the killing is worse. It is a worse state of affairs when someone dies (whether from being

killed or from natural causes) than it is when a person fails to come into existence, assuming that in both cases the lives would have been worth living.

In sum, it is difficult to see how Theory X can avoid implying that, other things being equal, (1) it is better to have a Happy Child rather than No Child; hence (2) there are serious moral objections to abortion, contraception, and celibacy; (3) the failure of a person to come into existence is at least as bad an outcome as the death of a person; hence (4) the failure to cause a person to exist is at least as bad as the failure to save a person's life and (5) the failure to cause a person to exist is at least as bad as killing a person. These claims, or at any rate the last three, are plainly unacceptable. The only hope for Theory X is that it can avoid having them as implications. Despite my earlier remarks, those who are attracted to the impersonal approach may remain optimistic. They may point out that there are, after all, some familiar candidates for Theory X that do not necessarily have these implications. If, for example, a Happy Child would have a level of well-being at or below the average, Average Consequentialism would not imply that it is better to have the Happy Child than to have No Child; yet it would imply that, *if* it were inevitable that *some* child was going to exist, it would be better to have the Happy Child than to have a Less Happy Child. But this is just an accident of the arithmetic. If the Happy Child would be above the average, it would be better, other things being equal, to have the Happy Child. And if the existing population were quite large and the Happy Child would be well above the average and would live long, it would be better, according to Average Consequentialism, to have the Happy Child than to save a person whose life was well below the average. It is important to note these facts, since Average Consequentialism is, in effect, concerned exclusively with the quality of life (which it measures in terms of the average) and gives no weight to increasing the number of lives except insofar as this affects the overall quality of life. Among the known impersonal theories of beneficence, therefore, it is the one least likely to have the claims cited above among its implications.[26]

IV CONCLUSION

To be acceptable, Theory X must imply that failing to save a person whose life would be worth living is, other things being equal, not just worse but significantly worse than failing to cause a person to exist. And this implication must not just be a contingent feature of the way the math works out. It must instead flow from the theory in a way that plausibly explains why the death of a person is a worse outcome than the failure of a person to come into existence.

[26] Average Consequentialism has been extensively criticized. See, e.g., 'Problems of Population Theory,' 111–15.

I cannot prove that no impersonal theory can satisfy this condition. Yet there is good reason to believe that no impersonal theory can. For it seems essential to the explanation of why the death of a person is worse than the failure of a person to come into existence that the former is worse *for someone* while the latter is not. Person-affecting considerations seem indispensable.

In many cases involving the Non-Identity Problem, a choice seems to have a bad effect but is nevertheless not worse for anyone. Parfit asks whether, in these cases, the fact that the choice is not worse for anyone makes a moral difference. 'There are,' he writes, 'three views. It might make all the difference, or some difference, or no difference. There might be no objection to our choice, or some objection, or the objection may be just as strong.'[27] Parfit accepts the third view, the No-Difference View. According to this view, whether or not the choice is worse for anyone is morally irrelevant; impersonal considerations alone matter. I have tried to show why I think this view will prove to be unacceptable. According to the first view, impersonal considerations have no weight; person-affecting considerations alone matter. As Parfit has shown, this first view is untenable. This leaves the second view.

As I understand it, the second view holds that an effect may be bad even if it is not worse for anyone, but not as bad as it would be if it *were* worse for someone. In short, impersonal considerations matter, but person-affecting considerations matter more. According to this view, the outcome in the Prenatal Case is worse than the outcome in the Preconception Case. In each case, an individual is caused to exist with a disability. In the Prenatal Case, this effect is worse for the individual, since he could have existed without the disability. In the Preconception Case, the effect is not worse for the individual; but it is worse impersonally, since a different child without the disability could have existed instead. According to Parfit's second view, the effect is worse when it is worse for the individual.

The three options cited by Parfit are not exhaustive. There is another view, which I will call the *Encompassing Account*, that I believe is more plausible than any of the three views Parfit mentions. It is similar to, but more complex than, the second view cited by Parfit. According to the Encompassing Account, person-affecting considerations and impersonal considerations are

[27] *Reasons and Persons*, 363.

distinct and nonadditive. Neither type of consideration is reducible to the other. Both matter; both provide reasons for action. But it is not always worse if a bad effect is worse for someone; sometimes it is, sometimes it is not.

Consider again the relevant effect in the Preconception and Prenatal Cases: a child is caused to exist with a disability. This effect may be worse for the child. When it is, it is *also* worse impersonally (for the world would have been better if this child had not been disabled). But the effect may be worse *only* impersonally, as in the Preconception Case. The Encompassing Account holds that, when this effect is worse for the child, as in the Prenatal Case, *that* fact provides whatever reasons there are to prevent it. That the effect is also worse impersonally is irrelevant. Yet, when the effect is worse only impersonally, as in the Preconception Case, that fact provides a reason to prevent it. In short, the objections to the Negligent Physician's action in the two cases are different. Contrary to Parfit, who claims that the relevant objection is impersonal in both cases, the Encompassing Account holds that the objection in the Preconception Case is impersonal while the objection in the Prenatal Case is that the effect is worse for the child.

How do the strengths of these objections compare? When an effect is both worse for someone and worse impersonally, as in the Prenatal Case, it is *at least* as bad as it would be if it were worse only impersonally. When an effect is worse only impersonally, as in the Preconception Case, it is *at most* as bad as it would be if it were bad for someone. Accordingly, the objection to the Negligent Physician's action in the Prenatal Case is at least as strong as the objection to his action in the Preconception Case. Thus it is possible, as Parfit claims and most of us intuitively sense, that the objections are equally strong in both cases. The impersonal objection in the Preconception Case may be as strong as the person-affecting objection in the Prenatal Case. It is not possible that the impersonal objection in the Preconception Case could be stronger, though it is possible that it could be weaker. To explain how the impersonal objection might be weaker, one might reason as follows:

> In the Prenatal Case, the disabled child could reasonably (though not in practice) have this thought: 'It could have been better for me.' That is a bitter reflection. In the Preconception Case, the parallel thought to which the disabled child would be entitled is: 'A better-off person might have existed instead of me.' This is not a disturbing thought; virtually all of us could reasonably believe this of ourselves. Thus, if we take up the points of view of the two children rather than surveying the possible outcomes from a distance, we have reason to think that the effect in the Prenatal Case is worse than that in the Preconception Case. This is not because the child in the Prenatal Case would actually have this thought and be made miserable by it. It is, rather, that the accessibility of this thought to the child reveals something important about the nature of the outcome.

Even if one accepts this reasoning, however, it is hard to believe that the moral difference between the two cases could be more than very slight.

As we have seen, there are other cases in which there seems to be a significant moral difference between an effect that is worse for someone and a corresponding effect that is worse only impersonally. Compare the failure to save a child who would otherwise have lived another seventy years with the failure to cause the existence of a child who would have lived for seventy years. In each case, we may suppose, there is a roughly equal loss: seventy years of life that would have been well worth living. The difference is that in the first case this loss is worse for the child—it is *his* loss—whereas in the second case the loss is worse only impersonally. The fact that, in the first case, the bad effect is worse *for someone* seems to make a significant difference. It seems that this must be part of the explanation of why the moral objection to failing to save the one child is significantly stronger than the objection to failing to cause the other child to exist.

According to the Encompassing Account, the morality of beneficence is governed by both impersonal and person-affecting considerations. In some instances, the two types of consideration may be of comparable strengths; in others, person-affecting considerations may be far stronger than corresponding impersonal considerations. This raises large questions. For any bad effect that is worse for someone, how do we determine what counts as the corresponding or 'equivalent' effect that is worse only impersonally? Why is the comparative strength of person-affecting considerations greater in some instances than in others? And how are the two types of consideration—person-affecting and impersonal—to be integrated into a unified account of our moral reasons? I cannot answer these questions. My aim here is more modest: to suggest how one can accept that the Impersonal Comparative Principle provides the correct account of the Preconception Case without committing oneself to the generalized No-Difference View—that is, the view that the whole of the morality of beneficence must be explained in impersonal terms. One can accept that there is a dimension to the Negligent Physician's conduct in the Preconception Case that can be criticized only in impersonal terms, that the Negligent Physician's conduct in this case is no less bad than his conduct in the Prenatal Case, but that his conduct in the Prenatal Case is objectionable for entirely different reasons. Thus it may be true that in the Prenatal Case the Negligent Physician owes compensation to the disabled child while this is not true in the Preconception Case.