

Caspar Hare
Published in *Ethics* 117, no. 3, 2007 – please cite the published version

December 2006

**VOICES FROM ANOTHER WORLD:
MUST WE RESPECT THE INTERESTS OF PEOPLE WHO DO NOT, AND WILL
NEVER, EXIST?**

Imagine that a couple pass up an opportunity to conceive and bear a child who would, they have every reason to think, be happy and healthy. Imagine that they have solid personal and financial reasons for doing so. If all other things are equal, many of us are inclined to think that they have done nothing wrong, and would be comfortable glossing this thought in the following kind of way: ‘Their decision was good for them, good for the over-burdened local school system... and bad for whom? Show me an *actual person* who has been made worse off.’

Now look back at your parents’ decision to bring you into the world. Unless you have had a very unpleasant life you will probably feel rather good about that decision, and will probably think poorly of anyone who wishes that it hadn’t been made. ‘After all,’ you might reasonably say, ‘that decision was good for an *actual person* – me.’

For those who are committed to a broadly consequentialist way of thinking about ethical questions, such considerations may lend some intuitive weight to *Moral Actualism*. Moral actualism says that the moral status of a token action is determined by its effects on the interests of all and only *actual* past, present and future people. The interests of merely possible people (those people who could have existed but didn’t, and those who might exist but won’t) are morally irrelevant.

A story may make the structure of moral actualism vivid. After the last of us has died, God prepares one list of all the people there have ever been, and another of all the

things they ever did. When he casts his mighty eye down the first list he thinks ‘I loved all these creatures individually, for all of their particular characteristics, and I loved them far more than I love any of the enormous variety of creatures that might have existed but didn’t.’ So when, on the Day of Judgment, he determines whether a particular action on the second list was right or wrong, he is only concerned with whether it made *the people on the first list* better or worse off. He looks to see if there was an alternative that would have been better for *these very people*, and cares not a bit about whether there were alternatives that would have been better for other people – people not on the list, people who never actually existed.

Moral actualism has attracted a good deal of attention and sympathy in recent years, in large part because it seems to follow quite naturally from two independently respectable doctrines: modal actualism and the so-called ‘person-affecting approach’ to ethics.

(Modal Actualism) There are no non-actual things.

(The PAA) The moral status of an action is determined, not by whether its outcome is simply better or worse than the outcomes of alternative actions, but by whether its outcome is better or worse for you, for me, for Joe Smith... etc. The morally relevant notion of betterness is *relational* – betterness for a particular person.

The person-affecting approach tells us that actions are made right or wrong by their effects on people. Modal actualism tells us that there are only actual people.

But, for all its surface appeal, moral actualism has been controversial, largely because it has been thought to have counter-intuitive practical implications. Consider:

*Childless George*¹

George passes up an opportunity to conceive a child who would, George has every reason to think, be relentlessly miserable from the start to the end of his or her sorry little life. By doing so, George imposes some small cost on actual people.

Hasn't George done the right thing, even though he has not benefited any actual people?

*The Two Inconsiderate Mothers*²

Two would-be mothers, June and Mary, knowingly cause their future children to have heart conditions, June by ignoring her doctor's advice and taking recreational drugs whilst pregnant, Mary by ignoring her doctor's advice and conceiving a child whilst recovering from German measles. June could easily have refused the drugs. Mary could easily have waited to conceive her child.

¹ Jan Narveson first drew attention to cases like this. See Narveson (1967).

² Derek Parfit first drew attention to cases like these. See Parfit (1976), and (1983) Chapter 16.

Don't *both* mothers do the wrong thing, even though June harms an actual person and Mary does not? (We may assume that Mary's actual child is better off existing than not, and that, if Mary had waited, her actual child would not have existed.)

Derek Parfit took cases like these to be deadly counter-examples to actualism. But he has not had the last word. In the *Childless George* case some philosophers insist that, strictly speaking, George has not done the right thing, but argue that there is an important sense in which he has been wise to avoid putting himself in the position of doing something very wrong.³ In the *Two Inconsiderate Mothers* case, some philosophers concede that Mary does wrong, but argue that her decision is in some morally relevant sense bad for her actual child.⁴ Others argue that, if the case really is as described, if it is true that Mary's actual child would not have existed had she waited three months, and nobody else has been made worse off by her conceiving immediately, then she has done nothing wrong – she is not a *baby machine*, whose job it is to make the best possible person.⁵ And perhaps it is not surprising that the cases have generated so much controversy, for when a moral intuition conflicts with a moral principle, the question of whether to discard the principle to preserve the intuition ('a deadly counter-example') or discard the intuition to preserve the principle ('a surprising moral discovery') is notoriously *delicate*. How one answers it often depends on how strongly one was inclined to believe in the principle in the first place.

In the first part of this paper I will argue that this debate has been, in an important sense, premature. In normative ethics, before we begin to consider whether a particular

³ See Parsons (2003)

⁴ See, for example, Woodward (1986)

⁵ See, for example, Roberts (1998) – a dogged and impressive effort to encourage us to reshape our intuitions about non-identity cases in line with actualism.

theory has plausible practical implications in common and fantastical cases, before we start worrying about whether it says intuitively correct things about run-away trolleys, fat men trapped in mine shafts, triage and such, we should ask whether the theory employs deontic concepts (like the concept of *obligation*, the concept of *agency*, the concept of a *range of accessible options*) in an acceptable way. Moral actualism fails on this count. If this has not been recognized it may be because the theory has never been precisely formulated. I will try to formulate it precisely, and show that it breaks down when applied to certain kinds of cases. My point, I should emphasize, is not that the would-be moral actualist is forced to say surprising, counter-intuitive things about rightness and wrongness in these cases. I will avoid appealing to any pre-theoretic moral convictions. It is that, *on pain of deontic absurdity*, a would-be moral actualist has literally no option but to say that in these cases significant aspects of the moral status of actual token actions are determined by how the alternatives would have affected people who never actually exist.

If we accept that actions are not always right (or wrong) in virtue of being better (or worse) for actual people, and we accept that there are only actual people, must we say that what matters is whether actions are better or worse simpliciter? I think not. In the second part of the paper I will sketch out a way of dealing with the problematic cases, which takes it that what matters is whether actions are in certain ways better with respect to the interests of people.

PART 1: Moral Actualism is False

1.1 Actualism

Let's begin by distinguishing a couple of ways to give precise content to the central insight of the moral actualist (henceforth just the *actualist*, but please don't confuse her with the modal actualist). Let a_1, \dots, a_n be actions available to an agent at a time and $a@$ be the one that is actually taken. Let $S@$ be the set of actual people (the set of people who actually exist, or have existed, or will exist), and, in general, let Sa_j be the set of people who would exist (in the past, present and future) if action a_j were to be taken. All actualists believe:

(Actualism) The moral status of $a@$ is determined by whether its outcome is better or worse *for people*⁶ in $S@$ than the outcomes of other available actions.⁷

What is the 'moral status' of an action, and how, exactly, does the determination work?

Theorists may answer this question in a variety of ways, but for present purposes all we

⁶ One should be wary, in this context, of talking about which outcomes are better *for a group of people*, since advocates of the person-affecting approach traditionally do not believe that conflicting interests can be aggregated in a straight-forward way (see Nozick (1974) Chapter 3, for example). If, for example, one outcome is slightly better for Hilary and Chelsea, and another is much better for Bill, it may be that there is no morally relevant sense in which one or the other is better *for the Clinton family*. So please take 'better for members of $S@$ ' as shorthand for 'better for actual person 1, better for actual person 2, better for actual person 3... etc.'

⁷ A sentence like 'I care about how things would have been for actual people if the agent had behaved differently' supports two readings. To use the idiom of worlds: on one reading, we take the term 'actual people' to pick out those people who exist in the speaker's world. (David Lewis called this 'the primary sense of 'actual''— see Lewis (1976) section IX). On another we go to the world in which the agent behaved differently, and then take the term 'actual people' to pick out the people who exist in that world. (Lewis called this 'the secondary sense of 'actual'''). As (Actualism) makes clear, the actualist's claim that the moral status of an act is determined, in part, by how things *would have been for actual people if the agent had behaved differently* should be read the first way. Read the second way, the theory would be vacuous – compatible with straight-forward act-utilitarianism, for example.

need to assume is that (Actualism) commits them all, at the very least, to one weak principle. Let Oa_1, \dots, Oa_n be the outcomes of a_1, \dots, a_n :

(Minimal Commitment) If, for some Oa_i in $\{Oa_1, \dots, Oa@, \dots, Oa_n\}$, $Oa@$ is no better than Oa_i for any member of $S@$, and $Oa@$ is significantly worse than Oa_i for some member of $S@$, and all other things are appropriately equal,⁸ then, given the range of alternatives $\{a_1, \dots, a@, \dots, a_n\}$, $a@$ ought not to be done.

If, for a given range of alternatives, an action is pareto-inferior to some alternative for the people who matter (no better for any of those people and significantly worse for some of those people) and all other things are appropriately equal, then, given that range of alternatives, it ought not to be done.

This is, as I say, a very weak commitment. Actualists will certainly want to develop a theory that covers more complicated cases, where the interests of actual people conflict. They will probably also want their theory to tell us whether actions ought to be done *simpliciter*, not merely whether they ought to be done *relative to a particular set of alternatives*. And they may think other deontic categories significant. They may want their theory to tell us whether actions are right or wrong, permissible or impermissible,

⁸This clause should exclude, for example, cases in which the relevant member of $S@$ deserves to be worse off, cases in which there is greater equality among members of $S@$ in $Oa@$ than Oa_i , and cases in which rights of members of $S@$ are violated in Oa_i but not $Oa@$. Versions of actualism may (for all we have said so far) attribute significance to whether actual people deserve what they get, to equality among actual people, and to whether the rights of actual people are violated in the various accessible worlds.

praise-worthy or blameworthy... etc. No matter. For present purposes all we need is the Minimal Commitment.

Actualism and the Minimal Commitment tell us something about the moral status of actual actions, but what about the moral status of non-actual actions? What if we want to determine whether it was (or is) morally permissible to do something that actually wasn't (or won't be) done? There seem to be two ways for actualists to go here. They could say that what matters is whether the outcome of the action would have been better for those people who would have existed had it been performed. Thus:

(Weak Actualism) The moral status of any a_j , actual or not, is determined by whether its outcome is better or worse *for people in Sa_j* than the outcomes of the other available actions.⁹

Or they could say that, even in these cases, what matters is whether the outcome of the action is better or worse for actual people. Thus:

(Strong Actualism) The moral status of any a_j , actual or not, is determined by whether its outcome is better or worse *for people in $S@$* than the outcomes of the other available actions.

⁹ Note that, for a modal actualist, when a_j is a non-actual action, talk of 'the people in Sa_j ' and of what outcomes are better or worse for them, is really shorthand for more complicated modal talk. So, for example, ' Oa_j is better than Oa_i for members of Oa_j ' is really shorthand for 'if a_j were taken, Oa_j would be better than Oa_i for members of Oa_j '. For economy's sake, I will use the shorthand here.

I suspect that most of the philosophers who are inclined towards actualism would, if confronted with this distinction, endorse weak actualism. Strong actualism is a prima facie implausible view. Consider this: we know what the strong actualist says about the moral status of non-actual a_j , but what does she say about what the moral status of a_j would be, if it were performed? She can say that what would matter would be whether it's outcome would be better or worse for people in $S@$. She is then committed to saying, for example, that if the population of the world had been entirely different then all actions would have been morally neutral – the genocidal adventures of non-actual Hitler would not have been wrong, because they would not have been bad for people in $S@$. That sounds implausible. Or she can say that what would matter would be whether its outcome would be better or worse for people in S_{a_j} . She is then committed to saying that there are circumstances in which the moral status of an action is different from what it would be if it were performed – e.g. there are circumstances in which it is permissible for me to pick up the pan, but if I were to pick up the pan then it would be impermissible for me to do so. That sounds implausible too, albeit in a different way.

But we need not dwell on the prima-facie implausibility of strong actualism, because both strong and weak forms of actualism are, for structural reasons, untenable.

1.2 Symmetrical Cases

I'll illustrate this by considering some cases in which an agent, call her Kate, has a range of actions available to her. In each case there's a set of people who exist whatever she does and a set of people who may or may not exist, depending on what she does. We can assume throughout that her decision has no substantive effect on the interests of

people in the former set (call it set $\cap S$) – none of the possible outcomes is better or worse for any of these people. And we can assume throughout (reasonably enough, I think) that this means the interests of people in $\cap S$ have no bearing on the moral status of the actions in question.

First, there are cases where Kate has exactly two actions available to her, each of which will lead to the existence of a different happy and healthy child. Call such cases *positively symmetrical*. Then there are cases in which each of the actions will lead to the existence of a different relentlessly miserable child. Call such cases *negatively symmetrical*. The argument that follows could be made about either kind of case. I'll focus on the latter kind.

A Negatively Symmetrical Case

Act	Outcome	People Who Exist
a_1	Baby Jack has a short life of relentless misery	$Sa_1 = \cap S + \{\text{Baby Jack}\}$
a_2	Baby Jane has a short life of relentless misery	$Sa_2 = \cap S + \{\text{Baby Jane}\}$

Kate must choose between bringing relentlessly miserable Jack, or relentlessly miserable Jane, into the world.

Confronted with such a case, one might be inclined to say: “Jack and Jane would both be worse off existing than not.” Let's assume, for the moment, that this is broadly

right. In this case the outcome of a_2 is better than the outcome of a_1 for Baby Jack, and the outcome of a_1 is better than the outcome of a_2 for Baby Jane.

1.3 Weak Actualism on Symmetrical Cases

It follows that, according to weak actualism, in this case *both* of the available actions are such that, given the alternatives, they ought not to be done – since the outcome of a_1 is in the relevant way worse for members of Sa_1 , and the outcome of a_2 is in the relevant way worse for members of Sa_2 . So Kate is *strongly fated*¹⁰ to do what she ought not to do, in that she cannot avoid doing what she ought not to do.

This shows that a weak actualist has odd ideas about obligation. Consider a deontic principle:

- (D) If, given alternatives $\{x,y\}$, x ought not to be done, then, given alternatives $\{x,y\}$, y ought to be done.

It seems very plausible indeed. After all, saying ‘Kate ought not to x ’ is really just a way of saying ‘Kate ought to not x ’ without splitting an infinitive. And that is just a way of saying ‘Kate ought to do something other than x ’. And if, given alternatives $\{x,y\}$, Kate ought to do something other than x , then surely, given those alternatives, she ought to do y .

¹⁰ The oracle tells me that I will kill my grandfather whatever I do. No matter whether I join the army or the priesthood, stay at home or go into exile, all of the branching paths before me lead to his death at my hands. Knowing the oracle to be right, I find myself in the grip of *strong fate*.

If the weak actualist were to accept (D) then she would be forced to say that, given alternatives $\{a_1, a_2\}$, both a_1 and a_2 both ought and ought not to be done. Perhaps this would not in itself be disastrous. Some philosophers, (though few consequentialists¹¹) are prepared to accept that there can be genuine moral dilemmas, and in these cases there are actions that are both right and wrong, actions that both ought and ought not to be done. I won't venture into this vast and well-trampled territory here. The important point for present purposes is just that a weak actualist cannot be one of these philosophers. If the weak actualist concedes that, given the alternatives, a_1 ought to be done, then she concedes that this important aspect of its moral status is determined by how alternative actions would have affected people who don't exist in the outcome of a_1 – a_1 ought to be done because a_2 ought not to be done, and a_2 ought not to be done because its outcome is worse for people in S_{a_2} , people who don't exist in the outcome of a_1 . But with this concession she has ceased to be an actualist at all.

So the weak actualist must deny principle (D). Can she do so? Well, it may be logically possible for her to do so, in the sense that there are substantive models of deontic logic that do not entail it. But this should lead us to be very suspicious of her notion of 'obligation'. Imagine that someone tells me, in some context, what I ought not to be doing:

“Given that moving and staying still are the alternatives available to you, you ought not to move”, says he.

“So, given those alternatives, I ought to stay still...”, say I.

¹¹ Standard consequentialism, whether expressed as a *maximizing* or *satisficing* theory, has it that at least one of the actions available to an agent (assuming there are finite number of them) is not wrong – one whose outcome is at least as good as any of the alternatives. For an extended discussion of this see Norcross (1995) and Carlson (1995).

“No”, says he, “you ought not to move, but it’s not the case that you ought to stay still.”

I would swiftly conclude that this person and I are talking past one another. Whatever he means by ‘ought’ is not what I mean by ‘ought’.

Nor do the weak actualist’s problems end here. Imagine that Kate actually creates miserable baby Jack in the negatively symmetrical case, and let’s pretend that we can make sense of what the weak actualist means when she says that Kate ought not to do it. There remains a clear sense in which the moral status of this action is influenced by the interests of non-actual Jane. To see this contrast the negatively symmetrical case with an asymmetrical case. In this case, if Kate had created Jane, Jane would have been perfectly happy:

An Asymmetrical Case

Act	Outcome	People Who Exist
$a_3(@)$	Baby Jack has a short life of relentless misery	$Sa_1 = \cap S + \{\text{Baby Jack}\}$
a_4	Baby Jane is thoroughly, happy and healthy	$Sa_2 = \cap S + \{\text{Baby Jane}\}$

In both cases, according to weak actualism, Kate actually does something that she ought not to do, but even the weak actualist must concede that there is an important difference between them. In the negatively symmetrical case moral considerations are deeply unhelpful to Jane. It’s not that morality doesn’t tell her anything – morality tells her not

to do both of the things she can do. It's rather that, on balance, moral considerations weigh no more strongly in favor of doing one thing than the other. But in the asymmetrical case moral considerations are helpful, they weigh more strongly in favor of one thing than the other (after all, a_4 ought to be done, a_3 ought not to be done). And surely this makes for a real difference in the moral status of a_1 and a_3 . a_3 is such that moral considerations weigh more strongly in favor of an alternative, but a_1 is such that moral considerations do not weigh more strongly in favor of an alternative. When, in the throes of deliberation, I wonder about the moral status of an action, this is precisely one of the things I am wondering about: do moral considerations weigh more strongly in favor of an alternative? And the aspiring weak actualist must concede that this real difference in the moral status of the respective actions results from the way that alternative actions would have affected non-actual people (after all, the cases only differ in this respect.)

1.4 Strong Actualism on Symmetrical Cases

A strong actualist (though committed, I remind you, to a prima facie implausible view) would appear to have an easier time addressing these cases. Strong actualism says that, in these cases, at most one action ought not to be done, the one whose outcome is worse for members of $S@$. So the strong actualist will not be forced to deny any compelling deontic principles.

But, for the strong actualist, morality remains deeply unhelpful in symmetrical cases. Imagine that you are Kate, a strong actualist, deliberating about what to do in the negatively symmetrical case. For the moment you haven't decided whether to do a_1 or a_2 ,

but you know, presumably, that you will do one or the other. Assume that you will actually a_1 . It follows that $S@ = Sa_1$, and, since the outcome of a_1 is worse for members of Sa_1 than the outcome of a_2 , you ought not to do a_1 and ought to do a_2 . Assume, on the other hand, that you will actually do a_2 . It follows that $S@ = Sa_2$ and, since the outcome of a_2 is worse for members of Sa_2 than the outcome of a_1 , you ought not to do a_2 and ought to do a_1 . You are, then, in the odd position of knowing, in advance of having made up your mind about what to do, that the action you will actually take is the one you ought not to take, and the action you could take but won't is the one you ought to take. You are *weakly fated*¹² to do what you ought not to do. It's not that you *can't avoid* doing what you ought not to do, it's just that you know you actually *won't*.

So what role might deontic facts play in your prospective deliberation about what to do in symmetrical cases? Well, they can't play any role at all – because you can't know which action you ought to take until *after* you have stopped deliberating and wholly committed yourself to one or the other. *Then* you know, but then it's too late. The fact that it ought not to be done is *deliberatively irrelevant*. Indeed, there is little point deliberating at all. As a moral deliberator your basic goal is to ensure that you actually avoid doing what you ought not to do, but if you are to achieve this goal in a negatively symmetrical case, you must actually do the thing that you do not actually do. And that, of course, is impossible.

¹² The oracle tells me that although, if I were to make the right choices, I would not end up killing my grandfather, it is a sad matter of fact that I actually will end up killing my grandfather. I am capable of taking any of the branching paths before me, only one of which leads to his death at my hands, but that is the one I will actually take. Knowing that the oracle is always right, I find myself in the grip of *weak fate*.

Now this observation¹³ is not *in itself* fatal for the strong actualist. He is free to say that his aim is to provide us with a theory about what makes it the case that acts have a certain moral status, not a theory that will serve as a guide to action in every case (in the jargon: he aims to give us a ‘criterion of objective moral rightness’ not a ‘moral decision procedure’). So a strong actualist need not be bothered by the fact that there are situations in which it is practically impossible to know the moral status of an action before it is performed, any more than (e.g.) an act-utilitarian need be bothered by the fact that there are situations in which it is practically impossible to know whether an act promotes utility before it is performed.

But compare, as before, your actually creating miserable Baby Jack in the symmetrical case with your actually creating miserable Baby Jack in the asymmetrical case. In the latter case the strong actualist must concede that the fact that your act ought not to be done *can* play a role in prospective moral deliberation. You could have thought like this: ‘Assume that I will actually do a_1 ; it follows that $S@=Sa_1$, and, since the outcome of a_2 is better for members of Sa_1 , a_1 is the act that ought not to be done. Now assume that I will actually do a_2 ; it follows that $S@=Sa_2$, and, since the outcome of a_1 is worse for members of Sa_2 , a_1 is again the act that ought not to be done. Either way I know that I ought not to do a_1 , so I’ll avoid doing it.’

Now the strong actualist has given a great deal away. There is an important difference between Kate’s actions in the negatively symmetrical and imbalanced cases. In the second case deliberatively relevant moral considerations weigh against her action, while in the first case they do not. Isn’t this a moral difference? The strong actualist must

¹³ A similar observation was made by Jeff McMahan in McMahan (1994). Thanks to him for bringing this to my attention.

either deny that it is or concede that aspects of the moral status of actual actions are determined, in part, by how the alternatives would have affected non-actual people.

1.5 The Moral Neutrality of Non-Existence

Is there any room for an actualist to maneuver here? In these next two sections I will briefly consider two lines of response.

First, I claimed that actualists cannot but concede that the interests of non-actual people bear upon the moral status of actual actions in certain kinds of cases – cases where, of a pair of outcomes, one is worse (or better) for people who exist in the outcome of the one and other is worse (better) for people who exist in the other. But this only cuts against actualism if it is at least conceptually possible that there be such cases. I assumed that it is. An actualist could deny this. An actualist could insist that people cannot be better or worse off existing than not. Call this view the *moral neutrality of non-existence*:

(MNNE) One outcome is better than another for a person only if that person exists in both outcomes.

MNNE is highly controversial.¹⁴ But that's beside the point. The point is that if the actualist accepts this controversial doctrine and thereby rules out the problematic two-option symmetrical cases, the familiar problems rearise in certain three-option cases.

Here is an example:

¹⁴ Many writers have considered or discussed it (see Parfit (1984), Appendix G, for example), but fewer have endorsed it. It is endorsed by John Broome in Broome (1999). A useful response is in Parsons (2002).

A Three-Option Symmetrical Case

Act	Outcome	Existing People
a_1	Baby Jack is healthy Baby Fred is unhealthy	$Sa_1 = \cap S + \{\text{Jack, Fred}\}$
a_2	Baby Fred is healthy Baby Jane is unhealthy	$Sa_2 = \cap S + \{\text{Fred, Jane}\}$
a_3	Baby Jane is healthy Baby Jack is unhealthy	$Sa_3 = \cap S + \{\text{Jane, Jack}\}$

In this case, assuming MNNE, weak actualism has it that, given the alternatives $\{a_1, a_2, a_3\}$, Kate ought not to do a_1 , a_2 or a_3 . The outcome of a_1 is worse than the outcome of a_2 for members of Sa_1 (it is worse for Baby Fred and neither better nor worse for the other members of Sa_1). The outcome of a_2 is worse than the outcome of a_3 for members of Sa_2 (it is worse for Baby Jane and neither better nor worse for the other members of Sa_2). And the outcome of a_3 is worse than the outcome of a_1 for members of Sa_3 (it is worse for Baby Jack and neither better nor worse for the other members of Sa_3 .)

And strong actualism again has it that Kate is in the odd position of knowing, before she acts, that she will actually do what she ought not to do, while she ought to do one of the things she could do but won't. If she actually does a_1 then she ought to do a_2 . If she actually does a_2 then she ought to do a_3 . If she actually does a_3 then she ought to do a_1 .

And this means, for just the reasons I gave in section 1.2 (I will spare you repetition of the details), that in these cases aspiring strong and weak actualists will end up committed to the moral significance of non-actual interests.

1.6 Perhaps Some Non-Actual People Don't Matter?

Second, I assumed that that an actualist must at the very least endorse:

(Minimal Commitment) If, for some Oa_i in $\{Oa_1, \dots, Oa_n\}$, $Oa@$ is no better than Oa_i for any member of $S@$, and $Oa@$ is significantly worse than Oa_i for some member of $S@$, and all other things are appropriately equal, then, given the range of alternatives $\{a_1, \dots, a_n\}$, $a@$ ought not to be done.

But perhaps actualists could deny this. Perhaps they could say that, although the moral status of an act is indeed determined by the interests of actual people, sometimes the interests of some actual people don't matter.¹⁵ In the negatively symmetrical case, for example, the fact that the outcome of a_1 is bad for actual Baby Jack has no bearing on the moral status of a_1 . In this case, although Baby Jack is an actual person, his interests don't count.

So when do the interests of an actual person count? One must tread carefully here. The easiest thing to say is that the interests of actual future people (people who don't yet exist, but will) don't count in just the problematic two- and three-option symmetrical cases. But this won't do, because it amounts to saying that in some cases the way that non-actual alternatives would affect non-actual people *does* influence the moral status of

¹⁵ Some actualists seem keen to leave this possibility open. They express the central actualist insight as 'an act is wrong only if there is some actual person for whom it is bad'. This does not entail that an act is wrong whenever there is some actual person for whom it is bad and no actual person for whom it is good. See, for example, Temkin (1987) and Roberts (1998). I have paraphrased both.

an action – for what makes the cases *symmetrical* is precisely the way that non-actual alternatives would affect non-actual people.

Perhaps one could say that the interests of actual people don't count in all those cases where they would not exist, were the agent to act in some other available way. So, for example, when I actually a_1 in the two-option negatively symmetrical case it doesn't matter that my decision is bad for Baby Jack, because he would not have existed had I a_2 -ed. And when I actually a_1 in the three-option symmetrical case it doesn't matter that my decision is bad for Baby Jack, because he would not have existed had I a_3 -ed.

This view has, to speak euphemistically, some *unsettling* practical consequences. Here's one: whenever someone is mulling over two options, which will affect the lives of future people in significant ways, if we present her with a third option that involves the immediate destruction of the world, then she is relieved of the need to consider how her actions affect future people. Here's another: even if a short life of relentless misery is worse for the person who lives it than no life at all, there is nothing wrong with creating such a life. But let's put them aside (after all, some philosophers are notoriously tolerant of unsettling practical consequences – Peter Singer once defended the plausibility of the latter by arguing that, although there is indeed nothing wrong with bringing a relentlessly miserable creature into being, it is wrong, having done so, not to kill it immediately¹⁶). The important point for present purposes is just that this view, like actualism, has *structural* problems when applied to symmetrical cases. Take the three-option symmetrical case. As a moral agent, when I consider the three available actions together I have no preferences, but when I consider them pair-wise I do have preferences (for doing

¹⁶ See Singer (1976). Thanks to him for bringing this response to my attention.

a_2 over doing a_1 , for doing a_3 over doing a_2 , and for doing a_1 over doing a_3), *intransitive* preferences. But this is at least prima-facie evidence that the view would commit us to a kind of practical irrationality.

The only viable remaining alternative seems to be to say that the interests of *no* actual future people count, no matter how near or far in the future they are, no matter whether they would or would not exist whatever I do. The moral status of a present action is determined by whether it is good or bad for actual people who do or have existed. This, at last, relieves an advocate of moral from structural problems, but at drastic cost.

Far better just to accept what should really have been obvious all along. Part of what makes my actions right or wrong is how the alternatives would have been for people who would have been around if I had taken them.

PART 2: A Different Person-Affecting Approach to Non-Identity Cases

2.1 Should We Appeal to Goodness Simpliciter?

Actualism is false. So what? What does this tell us about how to think about the moral status of an action that affects who ever exists?

It is natural to think that rejecting actualism involves accepting that the morally significant concept of betterness is not *betterness for a person*, but *betterness simpliciter*. If actions are not right or wrong to the extent that their outcomes are better or worse for people, then they must be right or wrong to the extent that their outcomes are better or worse simpliciter. So, if we are to make progress towards determining the rightness or wrongness of token actions in tricky non-identity cases, we must try to construct a theory

that tells us when, and in virtue of what, one outcome is simply better or worse than another.

But this is a mistake. If we reject actualism then all we are directly committed to saying is that certain kinds of modal facts have a bearing on the moral status of actions. Actualism has it that *some* modal facts matter in this way, namely:

(MF 1) Whether or not the alternative actions available to the agent would have been better or worse for actual people.

When we reject actualism we concede that other modal facts matter too, namely:

(MF 2) Whether or not the alternative actions available to the agent would have been better or worse for the people who would have existed had those actions been taken.

And we certainly can make this concession without committing ourselves to the idea that actions are right or wrong to the extent that their outcomes are better or worse simpliciter.

In the space that remains available to me, I will sketch out a way to approach the classic non-identity case, that of the two inconsiderate mothers, with this in mind.

The Two Inconsiderate Mothers

Act	Outcome	People Who Ever Exist
June, pregnant, takes drugs (@)	Baby Junette is born with a heart-condition	$Sa_1 = \cap S + \{\text{Junette}\}$
June doesn't take drugs	Baby Junette is born healthy	$Sa_2 = \cap S + \{\text{Junette}\}$

Act	Outcome	People Who Ever Exist
Mary conceives immediately (@)	Baby Mariette is born with a heart-condition	$Sa_1 = \cap S + \{\text{Mariette}\}$
Mary delays conception	Baby Other is born healthy	$Sa_2 = \cap S + \{\text{Other}\}$

Parfit and common sense say that, just as June does wrong by taking recreational drugs, so Mary does wrong by conceiving immediately. I agree. The challenge is to explain why.

Many philosophers have been inclined to argue that the wrongness of Mary's action stems from her having wronged *Mariette*. This requires some ingenuity. One strategy is to argue that Mariette is (despite appearances) worse off existing than not.¹⁷ Another is to argue that Mariette is worse off because she would have existed without a heart condition (and with a very different genetic make-up) if Mary had waited.¹⁸ Another is to argue that, although Mary has not made Mariette on balance worse off, Mary has, in an illegitimate way, imposed some harm upon her so as to impose some benefit upon her.¹⁹ Yet another is to argue that, although Mary has not in any way harmed Mariette, she has

¹⁷ See Benatar (1997)

¹⁸ You can adopt this line if you reject the view that Mariette was essentially born when she was born (or thereabouts) and essentially possesses the genes she possesses (or thereabouts). This is not as weird as it sounds. Lewisian counterpart theorists think that whether or not a sentence like 'if Mary had waited, Mariette would not have had a heart condition' is true depends upon which entities in nearby possible worlds are to be considered counterparts of Mariette, and that whether or not entities are to be considered counterparts is a context-sensitive matter. Such theorists are free to say that when we ask 'would Mariette have been born if Mary had waited?' the context induces a very *loose* counterpart relation, making 'yes' the appropriate answer. I have heard this suggested many times in discussion, but never seen it in print.

¹⁹ See Schiffrin (1999) and Steinbock and McClamrock (1994).

nonetheless wronged her, in the familiar way in which you can wrong people, by violating their rights, without harming them.²⁰

In light of the failure of actualism, these efforts seem unnecessary. In order to say that Mary does wrong we do not need to find some actual *victim*, wronged by Mary.

Other philosophers take a different tack. They argue that Mary does wrong because she has a certain kind of *impersonal* responsibility – a responsibility to nobody in particular. Some argue that this is a responsibility that any person has – e.g. to avoid bringing about suffering.²¹ Others argue that that it is a special kind of responsibility that only parents have – to avoid creating children whose lives will have certain features.²²

This seems to me broadly the right way to approach the problem. The challenge is to identify precisely the nature of Mary’s impersonal responsibility. I suggest that it is responsibility to avoid bringing about states of affairs that are in one particular way worse than other states of affairs – not *worse simpliciter*, but *de dicto worse for the health of her child*.

2.2 De Dicto Goodness

The best way to introduce the idea of *de dicto good for* is with a lame joke. Zsa Zsa Gabor is conversing with an interviewer...

Zsa Zsa: “Ah! People misunderstand me! They think that I am just a creature of leisure, that I do nothing useful, but they are wrong. I am constantly finding new ways to do good for people.”

Interviewer: “Like what?”

²⁰ See Woodward (1986)

²¹ See Harris (1998)

²² See Freeman (1997)

Zsa Zsa: “I have found a way of keeping my husband young and healthy, almost forever.”

Interviewer: “Eternal youth... that is quite a discovery! How do you do it?”

Zsa Zsa: “I get a new one every five years!”

What might Zsa Zsa mean? In what sense is her behavior good for the health of her husband? Well, her joke illustrates that there are two ways of understanding a locution of the form ‘better for the health of __’, where the ‘__’ is filled by a descriptive referring term (a term like ‘her husband’ or ‘the 35th President of the United States’):

(De re better): Where S1 and S2 are states of affairs, S1 is *de re better for the health of* __ than S2, when the thing that is actually __ is healthier in S1 than in S2.

(De dicto better): Where S1 and S2 are states of affairs, S1 is *de dicto better for the health of* __ than S2, when the thing that is __ in S1 is healthier in S1 than the thing that is __ in S2 is in S2.

I have made out this distinction in a rather technical way, but the central idea is quite simple and intuitive. What Zsa Zsa did, when she exchanged her old husband for a newer model, was de dicto good for the health of her husband – if she had stuck with her previous marriage her husband would have been getting long in the tooth, but as things are her husband is overflowing with youthful good health. But it wasn’t de re good for the health of her husband – her actual husband is no better off (vis a vis health, at least!) than her actual husband would have been if she had not married him.

So sometimes you can make things *de dicto* better in some way without making things *de re* better in that way. And likewise for making things *de dicto* worse. When I befriend someone suffering from cancer, what I do is *de dicto* bad for the health for my newest friend – if I had been unsociable my newest friend would have been thriving, but as things are my newest friend is sick. But it is not (I hope!) *de re* bad for the health of my newest friend – my actual newest friend is no less healthy for my befriending him.

2.3 The Significance of Some Kinds of De Dicto Betterness

When Mary ignores her doctor's advice, she makes things *de dicto* worse for the health of her child. Her actual child is less healthy than the thing that would have been her child would have been. Does this matter?

In *Reasons and Persons*, Parfit considers an effort to explain the wrongness of an action like Mary's by appeal to the observation that it is *de dicto* bad for her child (he does not use those words, but he acknowledges that 'in one sense the girl's decision was worse for her child', and he clearly has something like *de dicto worse for* in mind) but immediately rejects it, on the grounds that it 'does not appeal to a familiar moral principle.'²³ He then proposes his own explanation, which boils down to two claims. The first (largely implicit) is that, all other things being appropriately equal, it is wrong to make the world worse. The second is his 'same number quality claim':

²³ Parfit (1986) section 122, pp. 359-360.

‘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.’²⁴

There is something quite natural about this approach. One morally relevant way for things to be bad is for things to be de re bad for a person. Another is for things to be bad simpliciter. Since Mary has not made things de re worse for any actual person, a natural way to explain why she has done something wrong is to say that she has made things worse simpliciter.

But I say that the earlier sort of explanation remains on the table. By conceiving immediately, Mary makes things de dicto worse for the health of her future child, and this is something she should have been concerned to avoid.

Parfit rejected this kind of explanation because he thought that no explanatory work was being done by the observation that Mary makes things ‘worse’ in this strange way. Generally speaking, it does not reflect well or badly on people if they make things better or worse in this way – to use the example he gives, it does not reflect well on the skill of a general if he makes things de dicto better for his army by switching to the winning side in the heat of battle. Of course this is right. All of the examples above illustrate the same point. When I make things de dicto worse for my newest friend, though de re worse for no-one, I don’t do anything wrong. Should I avoid befriending the terminally ill for this reason? No. Likewise, when Zsa Zsa makes things de dicto better for her husband, though de re better for no-one, she doesn’t do anything right – indeed, the silliness of the thought that she thereby does something right is precisely what makes her joke (sort of a

²⁴ Ibid. page 360.

little bit) funny. And it is lucky that de dicto badness does not, generally speaking, matter because we are *always*²⁵ making things de dicto worse in some way or other. Right now I am facing my office door, and the person in the office to my left is healthier than the person in the office to my right. By choosing to face the door rather than the window behind me, I have made things de dicto worse for the nearest person to my right. Does that matter? Not at all.

But it does not follow from the fact that de dicto betterness is not *always* morally significant that it is *never* so. Sometimes it is appropriate to expect people to care about making things de dicto better in some ways. Here's an example:

*The Safety Officer*²⁶

Tess is a state safety officer, whose job it is to regulate those features of the automobile that protect its occupants in the event of a collision – air bags, crumple zones and so forth. Noticing that people in her state are not wearing safety belts, she implements some tough new regulations and, a year later, is pleased to discover evidence that they have been effective, that the severity of injuries sustained in automobile accidents has been reduced as a result of people belting up. She gives herself a pat on the back.

Suppose that a philosopher accosts Tess and says:

²⁵ Take the sentence-schema: 'I have made things de dicto worse for ___'. In any world in which there are two people, one healthy, one unhealthy, we can insert in place of the '___' a definite description that picks out an unhealthy person, but would have picked out a healthy person if I had acted differently, and the resulting sentence will come out true.

²⁶ Thanks to Mark Johnston for this example.

“Why are you acting so smug? What makes you think that you have been doing your job? Your job is to make things *better for the victims of automobile accidents*. But what you did made things much worse for the victims of last year’s automobile accidents. Accidents involve split-second timing. If you had just made it *illegal* to wear a safety belt, then most of those people would not have fumbled with the clip for five seconds before pulling out of their driveways, and, for most of them, the momentary, unhappy combination of conditions (e.g. the bicycle veering across the junction, the taxi-driver rubbing his weary eyes, the crates of olive oil tottering unsteadily on the flat-bed truck...) that led to the accident would never have arisen. Most of them would never been involved in accidents of any kind.”

How can Tess reply? She has done well, that’s clear, but *why*?

One might be tempted to answer that she has done well by making things collectively better, not for last year’s accident victims, but for some larger group of people – e.g. for all members of her state who were alive last year. After all, these people would have been around whatever she did (we may assume, rather artificially, that this is not a ‘non-identity’ case) and, collectively, they would have suffered more, nastier injuries in automobile accidents had the new regulations not been in place. Or one might be tempted to answer that she has done well by making things better simpliciter – a world with fewer, less nasty injuries is simply better than one with more, nastier ones.

Neither answer would be entirely satisfactory. It may be that Tess has made things better simpliciter or collectively better for the people in her state (whether or not she has is under-determined by my description of the case), but doing so is neither sufficient nor necessary for doing her job well. Imagine that she implements safety regulations whose

effect is to hasten the deaths of frail old people. It may be that she has made things simply better and better, collectively, for the people in her state (the old people die quickly and their children do not have to support them), but she has not done well. Now imagine that she implements safety regulations whose effect is to protect frail old people. It may be that she has made things simply worse and worse, collectively, for the people in her state (for converse reasons), but she *has* done well.

A satisfactory answer from Tess must acknowledge the fact that it is not her job to improve people's lives in any way she can. She would clearly be remiss if she were to spend her time attending to the distant needy or worrying about the consequences of her regulations for the medical profession – ‘hmm... fewer spinal injuries will mean less work for doctors...’ etc. It is her job to focus on the well-being of accident victims. Their well-being is the measure of her success or failure. So this is what she should say:

“Yes, what I did was *de re* worse for last year's accident victims. But we regulators always know, long before we decide what substantive measures to take, that the substantive measures we will actually take (whether they involve reducing speed limits a little, increasing them a lot, adding more traffic lights or spreading nails all over the roads of America) are going to be *de re* worse for accident victims. My job was not to make things *de re* better for last year's accident victims, but to make it the case that last year's accident victims were, collectively speaking, healthier than those people who would have been last year's accident victims would have been if I had acted otherwise. In brief, my job was to make things *de dicto* better for the health of last year's accident victims. And I did just that.”

There are numerous other examples. Here's one:

The Cancer-Researcher

Simon is a cancer-researcher. He develops a cancer-battling technology that, ten years later, comes to market with spectacular success, saving people from malignant forms of the disease that would previously have killed them.

Has Simon done well? Apparently so. But *not* in virtue of having made things de re better for cancer patients. Given that whether or not you get cancer is typically a chancy matter – having to do with whether a stray dose of cosmic radiation happens to strike the nucleus of one of your cells in such a way as to cause a very unusual kind of genetic mutation – it seems very likely that many of the patients who ‘benefited’ from Simon’s treatment would never have gotten cancer if it had not been invented.²⁷ Nor has he done well in virtue of having made the world better simpliciter, or collectively better for all members of his society. That is not his job. His job is to develop drugs to improve the condition of cancer patients. He has done well because actual cancer patients are better off than those people who would have been cancer patients would have been if he had not invented the treatment, which is to say that he has done well because he has made things de dicto better for cancer patients.

²⁷ The idea is that some people who get cancer are strongly disposed to do so (for genetic or environmental reasons) but others are just unlucky. For one of the unlucky people it will be true that (e.g.) if he hadn’t taken a walk at precisely such and such a time, and raised his neck to the sun at precisely such and such an angle, then most likely he would not have gotten cancer. And for any piece of technology that has had a major impact upon the lives of many people it will be true that if the technology had not been invented then most likely he would not have taken a walk at precisely such and such a time and raised his neck to the sun at precisely such and such an angle. You may not accept this, but for dialectical purposes that doesn’t really matter. The point is that we would not think worse of Simon and his invention if it turned out to be true – though it would then be the case that Simon’s invention has been de re bad for many cancer patients.

So there's a real psychological attitude that involves caring, not that the occupant of a certain role be as well off as possible, but that a certain role be filled by someone as well off as possible. And there are some situations in which we expect a person to exhibit this attitude, to care about making things *de dicto* better in some way. How could one argue that it is *appropriate* to have these expectations? I suggest that the following observations would suffice:

First, it is appropriate to expect the person in question to be *partial* – to have focused concern for Q, where Q is a person or group of people (the safety regulator should be concerned, not about the world or about people in general, but about accident victims, and the cancer researcher should be concerned, not about the world or about people in general, but about cancer patients.) And it is appropriate to expect concern for Q to play a significant role in guiding at least some of the choices she makes (at least some of the choices the safety regulator makes should be guided by her concern for accident victims; at least some of the choices the cancer researcher makes should be guided by his concern for cancer patients.)

Second, in virtue of the causal circumstances the person finds herself in, that partial concern has no *de re* expression. Whenever the person faces an array of options that have some impact upon the object of the concern, it is either the case that whatever she does will turn out to be *de re* better for Q, or that whatever she does will turn out to be *de re* worse for Q (*absolutely whatever* the safety regulator does in her role as a safety regulator will turn out to be *de re* worse for the accident victims upon whom she should focus, and *absolutely whatever* the cancer researcher does will turn out to be *de re* worse for the cancer patients upon whom he should focus). So *de re* concern for Q (the desire to

bring about states of affairs that are de re better for Q) cannot play a significant role in guiding any of the choices she makes. So, given that it is appropriate to expect concern for Q to play a significant role in guiding at least some of her choices, it is appropriate to expect her choices to be guided by de dicto concern for Q (the desire to bring about states of affairs that are de dicto better for Q).

The best argument for thinking that Mary has done something wrong in the classic non-identity case will run along similar lines. First is appropriate for Mary (as for any parent) to have a standing concern for the health of her children,²⁸ and it is appropriate to expect this concern to play a significant role in guiding at least some of those choices that affect the health of her children. Second, conceptions, like collisions, typically involve split-second timing – so much so that (unless we accept a very liberal view of trans-world identity) we seem forced to say that virtually all of the substantive decisions Mary makes before conceiving her child affect the identity of the child she conceives. This means that there is *no outlet* for de re concern for her child. As long as her future child will not be so very miserable that it would be better for it never to have existed, virtually *all* of the decisions she makes before conceiving it (whether they involve listening to her doctor's advice, ignoring her doctor's advice, starting up a college-fund or burning down the house) will turn out to be de re better for her future child. So, before her child is conceived, if concern for the health of her child is to play a role in guiding at least some of the choices

²⁸ The precise nature of this concern is controversial. Am I required to desire that things be better, *in every way*, for my children? One might think not. While I am obliged make sacrifices to ensure my child can walk, I am not obliged to make sacrifices so as to send her to an excellent private school when there are good public schools in our neighborhood. What is required of me (at least) is a narrower concern for my child's health. Am I required to desire that my child be *as healthy as possible*? One might think not. I am not obliged to feed her supplements that will turn her into an Olympian. What is required of me (at least) is a narrower concern that she be healthier in certain particular ways.

that affect her child (as it should) it must be expressed de dicto – as a concern that things be de dicto better for the health of her child.²⁹

2.3 Some Nice Features of this Approach

Mary does wrong by making things de dicto worse for the health of her child. Let me close by noting some advantages this approach to Mary's case has over the more familiar approach – which says that she does wrong by making the world simpliciter.

First, if we do wrong by making the world worse, Parfit's principle (Q) tells us that Mary does wrong if the same number of people will ever exist whatever she does – if this is a 'same-number non-identity case'. But, in the real world, very few cases like Mary's are same-number cases. Given that world-history is a large and encompassing thing, it seems likely that most decisions that affect who exists will reverberate through it for many generations, and unlikely that, when all is said and done, the numbers of people who ever exist will turn out the same whatever we decide. To take a very local, very small-scale case: I share one of my great grandmothers with forty six of my siblings, first cousins and second cousins. I think it unlikely that precisely this number of us would have been around if my great great grandparents had decided to delay her conception by a couple of months.³⁰

²⁹ Contrast Mary's case with an adoption case – Sarah makes things de dicto worse for the health of her child by adopting a blind child, Jeff, when there were sighted children at the orphanage. In Sarah's case, though de re concern for her child can play no role in guiding the particular decision about which child to adopt (assuming Sarah would substantially benefit any of the children by adopting them) it is not true that de re concern for the health of her child cannot play a role in guiding at least some of the choices she makes prior to the adoption. This is because not all of the significant things we do before adopting a child affect the identity of the child we adopt. Sarah could have made things de re worse for Jeff by giving away her family fortune before adopting him, by saying something cruel to him before adopting him...etc. Thanks to an anonymous reviewer for this example.

³⁰ This point (that 'same-number non-identity cases' are likely to be very rare indeed) is made by Jeff McMahan in McMahan (2005). Thanks to him for bringing it to my attention.

Given massive uncertainty about the long-term effects of actions like Mary's, we are entitled to apply (Q) to Mary's case only if we are entitled to make a simplifying assumption, and treat it as if it were a same-number case.³¹ On the de dicto betterness account no simplifying assumptions are called for. It doesn't matter whether Mary brings it about that more or less people ever exist. What matters is that she makes things de dicto worse for her child.

This is not to say that we never have to face different-number problems. What if a mother has to choose between conceiving twins or triplets? Just as, in the original case, it is appropriate to expect her to make things de dicto better for the health of her child, so in this case it is appropriate to expect her to make things de dicto better for the health of her children (a parent should care about the health of her children, but in these peculiar circumstances that concern has no de re expression.) To find out which outcome is de dicto better for her children we need to compare the collective health of the twins with the collective health of the triplets. I have no penetrating insights about how to go about making these sorts of comparisons. Sometimes they will be easy (e.g. severely unhealthy twins versus healthy triplets) but sometimes they will not (e.g. healthy twins versus asthmatic triplets, or asthmatic twins versus triplets, two of whom are healthy, one of whom has a heart condition). The familiar general strategies – *averaging*, *summing*... etc. – are available to us here, and the problems with these strategies are well known.

But the point remains that there are non-identity cases (cases like Mary's) about which the de dicto betterness account gives clear answers, though they may involve actions that affect how many people ever exist.

³¹ Whether or not we are entitled to make assumptions of this kind is a controversial matter – see Lenman (2000).

Second, the theory that Mary does wrong by making the world worse struggles to explain why her obligation to bring a healthier child into the world is so strong. After all, it is plausible to think that obligations to make the world better are not, generally speaking, just like obligations to particular people to make things better for them. Consider the following contrast, the sort of contrast of which critics of utilitarianism are fond:

Intra-Personal v. Inter-Personal Trade-Offs

Ned can, at no cost to himself or anybody else, either save his friend Alexis from enduring 19 hours of undeserved suffering or save his friend Adina from enduring 20 hours of undeserved suffering. He cannot save both. On a whim, a thoughtless impulse, he chooses to save Alexis. Adina suffers for 20 hours.

Rachel will suffer undeservedly. Her friend Bob can bring it about, at no cost to himself or anybody else, that she suffers for either 19 or 20 hours. On a whim, a thoughtless impulse, he brings it about that she suffers for 20 hours.

In each case someone brings it about that there are 20, rather than 19, hours of pain. But, says the critic, morally speaking, there is a world of difference between the cases. Bob senselessly imposes an extra hour of suffering upon Rachel. His action is bad for her and good for nobody – wrong by any measure. Ned’s action is not bad for somebody and good for nobody. It is good for Alexis. Hard-headed utilitarians may insist that he has done wrong by making the world worse simpliciter, but even they must surely concede that

Ned's obligation to make the world better simpliciter is less strong than Bob's obligation to make things better for Rachel.

In light of this, an advocate of the view that Mary does wrong by making the world worse is hard-pressed to explain why Mary's actions are, morally speaking, just like June's (June, recall, makes things *de re* worse for the health of her child by neglecting to take drugs mid-pregnancy.)³² June's action is bad for somebody and good for nobody. Mary's is not. Why doesn't that make a difference? Why, in this case, is Mary's obligation to make the world better simpliciter as strong as June's obligation to make things better for Junette?

We can explain the equivalence easily – both Mary and June fail to show appropriate concern for the health of their children. June should express that concern by making things *de re* better for her child. Mary should express that concern by making things *de dicto* better for her child. The Rachel and Ned cases would also be morally equivalent if it were appropriate to expect Ned to care about making things *de dicto better for his suffering friend*. But it is not appropriate to have that expectation. Perhaps Ned should have a standing concern for his friends, but this concern can be expressed *de re*. He can make things *de re* better or worse for (e.g.) Adina by choosing to save her or not.

Finally, you may have a feeling, as I do, that Mary's wrong-doing is in a certain way personal – in the way that there's a special kind of relationship between one who harms and

³² This is not to say that he has no options. For example, he could argue that *fairness* explains the difference between the Ned/Bob cases and the Mary/June cases. Ned behaves fairly. If he had chosen Alexis on the grounds that she would suffer one hour less pain this would have been unfair to her. But Mary does not behave fairly. If she had chosen to wait on the grounds that she would then conceive a healthier child, this would not have been unfair to Mariette (Mariette would not have existed.) But this turns on controversial ideas about fairness, ideas that someone sympathetic to the view that we do wrong by making the world worse simpliciter is unlikely to accept. Thanks to an anonymous reviewer for this observation. Or he could deny that the Ned and Bob cases are morally inequivalent. Or he could deny that Ned makes the world worse simpliciter by bringing about twenty, rather than nineteen, hours of suffering. But these are *uncomfortable* options.

one who is harmed, so there is a special kind of relationship between Mary and Mariette. Mariette has special grounds to feel aggrieved by what Mary did. But if Mary does wrong by making the world worse then this is hard to explain. After all, nobody has a special complaint against her. *Everybody* can complain ‘you have made things worse’. *Nobody* can complain ‘you have made things worse for me.’ The de dicto concern account puts a finger on Mariette’s special grievance. Mariette alone can say ‘you failed to show appropriate de dicto concern for your child, and *I am your child.*’

References:

- Benetar, D. (1997): “Why it is Better Never to Come into Existence” *American Philosophical Quarterly* 34
- Broome, J. (1999): *Ethics Out of Economics*, Cambridge: Cambridge University Press.
- Carlson, E. (1995): *Consequentialism Reconsidered*, Norwell MA: Kluwer Academic Publishers.
- Freeman, Michael (1997): *The Moral Status of Children: Essays on the Rights of Children*, The Hague, Kluwer Law International
- Lenman, J. (2000): “Consequentialism and Cluelessness”, *Philosophy and Public Affairs* 29
- McMahan, J. (1994): Review of David Heyd’s *Genethics: Moral Issues in the Creation of People*, *Philosophical Review* 103
- McMahan, J. (2005) “Preventing the Existence of People with Disabilities”, in David Wasserman, Robert Wachbroit, and Jerome Bickenbach, eds., *Quality of Life and Human Difference*, NY and Cambridge: Cambridge University Press.
- Narveson, J. (1967): “Utilitarianism and New Generations”, *Mind* 76
- Norcross, A. (1985): “Should Utilitarianism Accommodate Moral Dilemmas?”, *Philosophical Studies* 79
- Nozick, R. (1974): *Anarchy, State and Utopia*, New York: Basic Books.
- Parfit, D. (1976): “Rights, Interests and Possible People”, in Samuel Gorovitz ed., *Moral Problems in Medicine*, London: Prentice Hall
- Parfit, D. (1984): *Reasons and Persons*, Oxford: Oxford University Press
- Parsons, J. (2002): “Axiological Actualism”, *Australasian Journal of Philosophy* 80
- Roberts, M.A.(1998): *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law*, Oxford: Rowman and Littlefield.
- Roberts, M.A. (2002): “A New Way of Doing the Best That We Can: Person-Based Consequentialism and the Equality Problem”, *Ethics* 112
- Schiffirin, S. (1999): “Wrongful Life, Procreative Responsibility, and the Significance of Harm”, *Legal Theory* 5/2

- Singer, P. (1976): "A Utilitarian Population Principle" in Michael Bayles ed., *Ethics and Population*, Cambridge MA: Schenkman Publishing
- Steinbock, B. and McClamrock, R. (1994): "When is Birth Unfair to the Child?" *Hastings Center Report* 24/6
- Taurek, J. (1977): "Should the Numbers Count?", *Philosophy and Public Affairs* 6
- Temkin, L. (1987): "Intransitivity and the Mere Addition Paradox", *Philosophy and Public Affairs* 16
- Thomson, J. (2001): *Goodness and Advice*, Princeton: Princeton University Press
- Woodward, J. (1986): 'The Non-Identity Problem', *Ethics* 96