

Solving the Non-Identity Problem: A Reply to Gardner, Kumar, Malek, Mulgan, Roberts and Wasserman

DAVID BOONIN

University of Colorado Boulder

ABSTRACT

The non-identity problem arises from the fact that five premises that seem plausible yield a conclusion that seems implausible. Most proposed solutions to this problem rest on rejecting one of the argument's premises. In *The Non-Identity Problem and the Ethics of Future People* (2014), I argue against such solutions and propose that accepting the argument's conclusion ultimately provides a more satisfactory response to the problem. In this paper, I reply to some of the main claims made in response to the arguments in my book by the contributors to this symposium. In section 1, I respond to Janet Malek's critique of my discussion of the argument's first premise. In section 2, I consider Molly Gardner's proposal that we reject the argument's second premise in favor of what she calls the existence account of harm, identify some problems with that account, and argue that even if we accept it, this would fall short of solving the non-identity problem. In section 3, I respond to Rahul Kumar's rights-based solution, which rests on rejecting the fourth premise of the non-identity argument, by identifying a trilemma for that solution and arguing that none of the three options it identifies offers a promising route to a successful solution. In section 4, I comment on David Wasserman's discussion of three arguments that might be thought to point to a satisfactory solution by rejecting the fourth or fifth premise of the argument. In section 5, I consider and respond to Tim Mulgan's rule-consequentialist solution, which involves rejecting the fifth premise of the non-identity argument. And in section 6, I respond to Melinda Roberts' claim that even if my approach to the non-identity problem is successful in two-option cases, it cannot be successfully extended to three-option cases.

Keywords: non-identity problem, future generations, procreative ethics, harm, rights.

1. INTRODUCTION

I am grateful to all the contributors for their careful and challenging responses to the arguments in my book. While I can't address everything they say here, I will try to offer some thoughts about at least some of the main points that each of them raises. I will begin by considering whether my treatment of the case of Wilma can survive the scrutiny aimed at my discussions of P1 (Malek), P2 (Gardner), P4 (Kumar, Wasserman), and P5 (Mulgan, Wasserman) of the non-identity argument and conclude by considering whether, even if my treatment of the case of Wilma is successful, there is a problem with my extending it to the case of Risky Policy (Roberts).

1. CHALLENGING P1: A REPLY TO MALEK

P1 of the non-identity argument maintains that Wilma's act of conceiving now rather than taking a pill once a day for two months before conceiving does not make Pebbles worse off than she would otherwise have been. Janet Malek suggests that "Breaking P1 into two sub-premises can . . . offer conceptual clarity about the relevant objections and rebuttals" by distinguishing between those objections that involve challenging the first sub-premise and those that involve challenging the second. The two sub-premises she identifies are (1) The Non-Identity Premise, which says that if Wilma had conceived two months later than she actually did, a different child would have been brought into being and Pebbles never would have existed, and (2) The Life Worth Living Premise, which says that Pebbles is not made worse off than she otherwise would have been by Wilma's decision to conceive if her life is worth living (2019: 18). Strictly speaking, P1 does not follow from these two "sub-premises". What follows is the conditional claim that if Pebbles has a life that is worth living and if having a life worth living is not worse than never existing, then P1 is true. As a result, the sense in which Malek's typology amounts to breaking P1 into two sub-premises is not entirely clear. Her distinction may nonetheless prove useful, though, so I will follow it here in organizing my response to her discussion of the five objections to P1 that I consider in chapter 2 of my book.

Malek relates three of these objections to the Life Worth Living Premise:

the Incoherence Objection, the Equivocation Objection, and the Asymmetry Objection. She characterizes my discussion of these three objections as “relatively convincing in demonstrating that these objections are not sufficient to solve the non-identity problem” (2019: 19) and says that my “arguments about their limitations are correct” (2019: 21). Since she doesn’t try to defend any of these objections to P1, I won’t say anything against them here. But Malek also says that “Boonin’s analyses of these three objections reveal a fundamental conceptual flaw in his account” (2019: 21), and I will try to say something about that concern. The heart of Malek’s complaint here seems to be that these three objections “address the question of whether a life is worth living, therefore comparing existence and non-existence, which is *not* the question at the heart of the non-identity problem as described in Wilma’s case. Boonin therefore tackles three arguments that are conceptually irrelevant to the problem he wants to solve” (2019: 21).

I find this claim about my account unconvincing because I don’t think we can know what question lies at the heart of the non-identity problem in the case of Wilma until we settle on a solution. Consider, for example, the solution that Gardner defends elsewhere in this symposium (2019). On that account, Wilma’s act harms Pebbles even if it’s true that Wilma’s act doesn’t make Pebbles worse off than she would otherwise have been. If Gardner’s argument is successful, then whether Pebbles is the same person as Rocks is irrelevant to solving the problem and heart of the problem lies in our failure to correctly understand what it is to harm someone. Or suppose Kumar is correct, and that Wilma’s act violates someone’s rights even if it doesn’t harm anyone (2019). If Kumar’s argument is successful, then whether Wilma harms Pebbles is irrelevant to solving the problem regardless of whether Pebbles and Rocks are the same person, and the heart of the problem lies in our failure to see the many ways in which an act can wrong someone without harming them. Similarly, it seems to me, one cannot say in advance whether any of the three objections to P1 that Malek dismisses as irrelevant will succeed in explaining the wrongness of Wilma’s act. One has to critically evaluate them first. Perhaps it will turn out, for example, that it’s incoherent to say that Wilma’s act does not make Pebbles worse off than she would otherwise have been but not incoherent to say that Wilma’s act does make Pebbles worse off than she would otherwise have been. Or maybe the Asymmetry Objection can show in a satisfying way that Wilma would harm Pebbles more by conceiving her than she would harm Rocks by conceiving him and that this would suffice to show that it was wrong for her to conceive the former child rather than the latter. If arguments such as these turn out to succeed in showing that Wilma’s act is wrong, then it will turn out that the relationship between existing and

not existing does indeed lie at the heart of the non-identity problem in the case of Wilma. So, I agree with Malek that these objections are ultimately unsuccessful, but I'm not convinced that the task of critically evaluating them is irrelevant to solving the non-identity problem in the case of Wilma.

While I disagree with Malek about the relevance of the first three objections, I do agree that they are not the strongest objections to P1. The most promising, I think, is one of the two she relates to the Non-Identity Premise, the *De Re/De Dicto* Objection as developed by Caspar Hare (2007). The other objection she discusses in that context, the Metaphysical Objection, strikes me as considerably less significant. This is not for the reasons I tentatively endorse in the book and that Malek discusses in her paper, but because of what I go on in the book to refer to as the "more fundamental problem" with the objection (2014: 41-42), a problem that Malek does not seem to address at all. The problem is that virtually everyone who thinks that Wilma's act is wrong thinks it's wrong even if they don't endorse (or even know about) the kind of metaphysical picture such as David Lewis's on which it turns out that "Pebbles" and "Rocks" literally refer to one and the same person. What they think, as I put it in the book, is that "it is wrong for Wilma to conceive a blind child now rather than a sighted child later even if we assume that the sighted child really would be a different person" (2014: 41). And since the Metaphysical Objection claims that the sighted child would not be a different person, it can clearly do nothing to vindicate this conditional judgment.

The *De Re/De Dicto* Objection is much more promising in this respect because it presupposes that Pebbles and Rocks are two distinct people. I therefore turn to Malek's critique of my responses to Hare. Hare's argument is grounded in the case of Tess, whose job involves safety regulations relating to automobile collisions (2007: 516). Tess implements some new regulations and one result is that the severity of injuries suffered in car accidents is significantly reduced. But another result is that, because the regulations affect various decisions that drivers make, different accidents occur than would have occurred had the regulations not been implemented, and different drivers are injured as a result. Because of Tess's regulations, for example, when Jill is in an accident, she suffers a relatively minor injury, but if the regulations had not been implemented, Jill would not have been in an accident at all and Jack would instead have been in an accident, which would have caused him a more serious injury than the one that Jill suffered. Now suppose after Jill's accident, we ask whether Tess's regulations made things better for the "accident victim". Well, they certainly didn't make things better for the "accident victim" in the *de re* sense of that term. The accident victim in the *de re* sense is Jill and Jill is worse off because of the

regulations. Had the regulations not been implemented, Jill wouldn't have been in an accident in the first place. But the regulations did make things better for the "accident victim" in the *de dicto* sense of that term: the injury suffered by "whoever ends up in an accident" was less severe than the injury suffered by "whoever ends up in an accident" would have been had the regulations not been implemented. Hare counts on the reader to agree that, in the case of Tess, her obligation is to accident victims in the *de dicto* sense, not in the *de re* sense, and this seems quite plausible.

But, of course, the case of Tess can help us solve the non-identity problem only if it turns out to be relevantly similar to the case of Wilma. Hare argues that it is. His argument rests on two claims (2007: 519). First, the reason that Tess has obligations to accident victims in the *de dicto* sense is that (a) Tess has a special role that gives her special duties to a particular class of people, namely "accident victims," and (b) because which particular people will end up as "accident victims" depends on what choice she makes, she can't help "accident victims" in the *de re* sense; she can only help them in the *de dicto* sense. Second, both of these features of the Tess case that explain why Tess has *de dicto* obligations apply to the case of Wilma: (a) Wilma has a special role that gives her special duties to "her child," and (b) because which particular person will end up as "her child" depends on what choice she makes, she can't help "her child" in the *de re* sense; she can only help "her child" in the *de dicto* sense. If Hare is right that the best explanation of why Tess has *de dicto* obligations entails that Wilma does too, then we can solve the non-identity problem by saying that Wilma's choice made "her child" worse off than "her child" would otherwise have been in the *de dicto* sense.

In my book, I raise two main objections to Hare's solution (2014: 34-8). First, I argue that the *de dicto* obligation that Tess has is best understood as a duty to minimize the harm that accrues to whoever turns out to be an accident victim rather than to maximize the well-being of whoever turns out to be an accident victim: the regulations are justified because as a result of implementing them, Jill suffers less harm than Jack would have suffered. If that's right, then even if Hare is correct that Wilma has the same kind of duties that Tess has, this won't solve the problem because no matter which child Wilma conceives, that child won't be harmed by being conceived. Second, I argue that Hare is mistaken in claiming that Tess has the obligations she has in virtue of (a) and (b) noted above and argue instead that she has the obligations she has because she explicitly consented to take them on when she agreed to take her job. If that's right, then it doesn't matter what kind of obligations Tess has. We can't move from the claim that Tess has *de dicto* obligations to accident victims to the

claim that Wilma has parallel *de dicto* obligations to her child because unlike Tess, Wilma never explicitly agreed to take on such obligations.

In pressing my first objection to Hare's argument, I point out that if we thought Tess's obligation really was to make the outcomes better for accident victims *de dicto*, rather than to minimize the harm they incur, then we'd have to think that Tess should try to redirect accidents from people who are less healthy to people who are more healthy, even if the more healthy people would suffer larger losses than would the less healthy people, just as long as the healthier people would still be better off after their more serious accidents than the less healthy people would have been after the less serious accidents they would have had. Assuming that implication is unacceptable, we should agree that Tess's obligations are really to minimize harm to whoever ends up in an accident, not to maximize the well-being of whoever ends up in an accident. Malek objects to my *reductio* against Hare here by saying that there is an important difference between the case of Tess and the case of Wilma (2019: 23). But whether there is such a difference is irrelevant to my use of the example. The point of my example is to show that the kind of obligation that Tess has (minimizing harm to accident victims, not maximizing the well-being of accident victims) is one such that, even if Hare is right to say that Wilma has the same kind of obligation that Tess has, cannot show that Wilma's act is wrong (because Wilma's act does not fail to minimize harm to her child). My example shows this regardless of whether there are differences between Tess and Wilma. And Hare is the one arguing that the case of Wilma is analogous to the case of Tess, not me. So Malek's point that there are differences between Tess and Wilma makes things worse for Hare, not for me.

Malek also suggests that my response to Hare here in part depends on circular reasoning (2019: 23): I argue that if I'm correct that Tess's obligation is really to minimize harm to whoever ends up being in an accident, then Wilma's act doesn't violate that obligation even if Hare is correct that Wilma has the same kind of obligation that Tess has. After all, if Wilma conceives Pebbles she doesn't harm Pebbles and if she conceives Rocks, she doesn't harm Rocks. So she can minimize the harm that she causes to whoever ends up being her child by conceiving either child. Malek objects that this part of my argument "only goes through if its conclusion is assumed: that Pebbles and Rocks should be referred to as two different possible children" (2019: 24). I certainly agree that my argument here depends on the assumption that Pebbles and Rocks are two different possible people. But I see nothing circular or objectionable about assuming this. If we thought that Pebbles and Rocks were literally one and the same

person, after all, there would be no need to appeal to *de dicto* harm in the first place. We would simply say that Wilma's act makes that particular person worse off than that very same particular person would otherwise have been. The entire point of Hare's argument is that in the case of Tess, there really are different people who might end up as the accident victims, and that Tess has an obligation to "the accident victim" in the *de dicto* sense, regardless of which of those particular people end up fitting that description. Hare agrees that Pebbles and Rocks are accurately described as two different, distinct possible children, and he wants to argue that Wilma has a *de dicto* obligation to "her child" in the *de dicto* sense, regardless of which of those particular people ends up fitting that description, because he thinks this follows from what he has to say about what he takes to be the parallel case of Tess. In following Hare in assuming that Pebbles and Rocks are different possible people, then, I'm simply meeting Hare on his own terms and discussing his objection in the only context in which it can sensibly arise.

Finally, in pressing my second objection to Hare, I in part point out that if Hare were correct that the best explanation of why Tess has the obligation she has is because of (a) and (b) noted above, then in deciding which people to take on as patients in my new medical practice, I should make the choice that would be better for "my patients" in the *de dicto* sense, and so select healthier people to have as patients, because in the case of a doctor who does not yet have particular patients, the parallel versions of (a) and (b) are also true: as a doctor I have special obligations to "my patients" and because which particular people will end up as "my patients" depends on what choice I make, I can't help "my patients" in the *de re* sense; I can only help "my patients" in the *de dicto* sense. Assuming the implication that I should therefore choose healthier people to have as my patients is unacceptable, we should agree that Tess's obligations are best explained in some other way, and I go on to suggest that the best explanation is the simplest one: Tess explicitly agreed to take on those obligations when she took the job. Malek again objects to my *reductio* against Hare by saying that there is an important difference between the case of Tess and the case of Wilma (2019: 23). But whether that is so is again irrelevant to my use of the example. The point of my example is to show that Hare's explanation of the source of Tess's obligations has unacceptable implications. This has nothing to do with whether Hare is justified in trying to move from the case of Tess to the case of Wilma. And, again, it is Hare who is claiming that the case of Wilma is relevantly similar to the case of Tess, not me. So Malek's point that there are differences between the two cases again makes things worse for Hare, not for me. Regardless of whether Malek is right to suggest that some of the objections to P1 that I consider in my book are literally

irrelevant to solving the non-identity problem, then, I continue to believe that none of the objections to P1 are strong enough to solve it.

2. CHALLENGING P2: A REPLY TO GARDNER

Molly Gardner's important work on the existence account of harming (2015) was published after I finished writing my book, so I am particularly grateful for the opportunity to engage with it here. In several respects, Gardner's account represents a significant advance over other harm-based approaches to the non-identity problem, but I will focus here on explaining why I think it nonetheless falls short of providing a satisfactory solution.

Gardner formulates the existence account of harming as follows (2019: 34):

Harming (def.): An event, E, harms an individual, S, if and only if E causes a state of affairs that is a harm for S.

Harm (def.): A state of affairs, T, is a harm for an individual, S, if and only if

- (i) There is an essential component of T that is a condition with respect to which S can be intrinsically better or worse off; and
- (ii) If S existed and T had not obtained, then S would be better off with respect to that condition.

As Gardner explains, “the account asks us to compare the world containing both the victim and the alleged harm to a world where the victim exists without the alleged harm; we then check to see whether, in the latter world, the victim is better off” (2019: 34). If the answer is yes, then the event harmed the victim.

The existence account of harming seems to solve the non-identity problem by showing that Wilma's act of conceiving Pebbles really does harm Pebbles. If we compare a world where Pebbles exists and is blind to a world where Pebbles exists and is not blind, we find (I am assuming) that in the latter world, Pebbles is better off. This establishes that the state of affairs in which Pebbles is blind is a harm for Pebbles. Wilma's act caused that state of affairs, and so Wilma's act caused a harm for Pebbles. An event harms someone if it causes a harm for them. And so, according to the existence account, Wilma's act harmed Pebbles. Gardner maintains that this solution satisfies both the robustness requirement and the modesty requirement, and since it seems clear that it also satisfies the independence requirement, this would mean that it successfully solves the problem.

A reason for rejecting a premise of the non-identity argument satisfies the robustness requirement if it is strong enough to warrant rejecting any weakened version of the premise that would still be strong enough to generate the conclusion that Wilma's act is not wrong. I do not think Gardner's solution satisfies this requirement. This is because if we replace P2 with the existence account of harming, we can simply revise P4 in a way that will leave it just as plausible as the original version of P4 and that will still suffice to entail that Wilma's act is not wrong. Even if Gardner's account can help to show that Wilma's act harms Pebbles, that is, it cannot help to show that Wilma's act wrongs Pebbles, and so cannot help to solve the non-identity problem.

The reason for this is as follows. If P2 is true, it's easy to account for the moral relevance of harm. Suppose, for example, that Pebbles was not born blind and that Wilma harmed Pebbles in the P2 sense of harm by poking Pebbles' eyes out shortly after Pebbles was born, thereby making Pebbles worse off than she would have been had Wilma not poked her eyes out. Clearly Pebbles would later have good reason to feel resentment about Wilma's act, to feel disrespected by it, to wish Wilma hadn't done it, to feel entitled to an apology and compensation from Wilma, to think others would have been justified in dissuading Wilma from doing it, and so on. But if the existence account of harm is true, then the mere fact that A has harmed B does not, in itself, seem to be morally relevant. In the case where Wilma conceives Pebbles, for example, the existence account justifies Pebbles in saying something like this to Wilma: "your act harmed me because even though it produced the best result for me that you could possibly have brought about, it nonetheless made things worse for me than they would be in an alternative world that you could not possibly have brought about, one in which I exist and am not blind." To me, at least, it seems clear that this fact about Wilma's act does not give Pebbles a good reason to feel resentment about Wilma's act, to feel disrespected by it, to wish Wilma hadn't done it, to feel entitled to an apology or compensation from Wilma, to think others would have been justified in dissuading Wilma from doing it, or to have any other response we would think justified in a case where Wilma had wronged Pebbles.

If this is correct, then we can simply accept the existence account of harm in place of P2 and replace P4, which says that if an act does not harm anyone, then the act does not wrong anyone, with P4*: if an act does not harm anyone by making them worse off than they would otherwise have been, then the act does not wrong anyone. The result would then be that Wilma's act does harm Pebbles, but not in a way that could make her act wrong. In order for the existence account argument against P2 to satisfy

the robustness requirement, then, it would have to provide a reason for rejecting P4* -- a reason, that is, for thinking that harming a person in the existence sense of harm is the kind of thing that can wrong a person even in cases where the act does not make that person worse off than they would otherwise have been and even in cases where the act leaves them as well off as the person who did the act could possibly have left them. I have difficulty seeing what kind of reason that would be, and Gardner does not provide one. She provides reasons to think Wilma's act harms Pebbles but no explanation of why harming Pebbles in this sense would make Wilma's act wrong.

A solution to the non-identity problem fails to satisfy the modesty requirement if it has implications that are less plausible than the claim that Wilma's act is not wrong. Gardner identifies two somewhat abstract implications of her solution that might be thought to pose problems in this respect and argues that these implications are not that implausible (2019: 37). I have nothing to say about her discussion there, but I do think the existence account of harm has more specific and concrete implications that render her solution unable to satisfy the modesty requirement.

Consider, for example, the case of Mort. Mort is mortal. If he weren't mortal, he would be immortal. Let's assume that would be better for him. In that case, the existence account of harm entails that Mort's parents harmed him by conceiving him. That account, remember, "asks us to compare the world containing both the victim and the alleged harm to a world where the victim exists without the alleged harm; we then check to see whether, in the latter world, the victim is better off." If the answer is yes, then the event harmed the victim. If we compare a world where Mort exists and is mortal to a world where Mort exists and is not mortal, we find that in the latter world, Mort is better off. This establishes that the state of affairs in which Mort is mortal is a harm for Mort. Mort's parents' act of conceiving Mort caused this state of affairs, and so their act caused a harm for Mort. An event harms someone if it causes a harm for them. And so, according to the existence account, Mort's parents harmed Mort by conceiving him. But, of course, all of us are mortal. So the existence account entails that every act of conception harms the person who is conceived. This strikes me as much less plausible than the claim that Wilma's act is not wrong. In addition, if harming Pebbles in the existence sense of harm really were enough to make it wrong for Wilma to conceive Pebbles, as would be necessary for the account to satisfy the robustness requirement, we would have to say not just that every act of conception harms the person who is conceived, but that every act of conception is wrong. That is an even stronger reason to think the solution fails the modesty requirement.

And Mort is just the tip of the iceberg. Candice was born lacking an immunity to cancer. Xavier was born without X-ray vision. Mindy can't read other people's minds. Assuming all of them would be better off if they existed and did not have these conditions, the existence account entails that their parents harmed them (and, if the account is to satisfy the robustness requirement, did something wrong) by conceiving them. And these cases would also generalize to all acts of conception. Perhaps an argument could be given for the claim that lacking vision is a "condition" while lacking X-ray vision (or being mortal) is not. But I find it hard to picture what such an argument would look like and even harder to see why appealing to it would not be ad hoc. If there is a sense in which lacking vision is a "condition" but lacking X-ray vision is something else, why would that distinction be relevant to the question of whether the state of affairs in which the person existed in that state was a harm to that person? Of course, we could point out that it was impossible for their parents to conceive Mort, Candice, Xavier, and Mindy without these conditions. But it was equally impossible for Wilma to conceive Pebbles without Pebbles being blind. And if a defender of the existence account of harm were to bite the bullet and agree that everyone is harmed by being conceived, this would just reinforce the concern that harm in this sense is morally irrelevant.

A second kind of problem arises if we consider the case of Betty. Betty has a slight preference for having a child and knows that if she ever conceives a child, the child will be incurably blind. She goes ahead and conceives Bamm-Bamm and, as a result, Bamm-Bamm is incurably blind. The existence account of harming entails that Betty harms Bamm-Bamm by conceiving him for the same reason it entails that Wilma harms Pebbles by conceiving her. If harming Pebbles in this sense really does wrong Pebbles, as must be the case if the existence account is to satisfy the robustness requirement, then harming Bamm-Bamm in this sense must wrong Bamm-Bamm. But the claim that it's wrong for Betty to conceive a blind child when any child she conceives will be incurably blind strikes me as far less plausible than the claim that it isn't wrong for Wilma to conceive Pebbles. This provides a second kind of reason to think that a solution based on the existence account of harming will fail to satisfy the modesty requirement.

Of course, as Gardner notes toward the end of her discussion, if we accept P2 and reject the existence account of harming, we will have to accept other implications that may themselves seem implausible. Gardner focuses, in particular, on P2's implication that Shooter does not harm Victim in *Preemption* and that Angry Alastair does not harm you when he slaps you in the face (2019: 42). I agree that these implications initially

seem implausible, but they strike me as considerably less implausible than the implications the existence account has in cases like *Mort* or *Betty*. So even if we look at things just in terms of initial plausibility, I think P2 already comes out well ahead. In addition, I think the strategies I deploy in the book to show that these implications of P2 are not really as implausible as they may at first seem can survive the challenges that Gardner provides here.

Suppose you are wondering, for example, whether your intuition in the *Trolley* case really is, as I assume, that pulling the switch doesn't harm Philip or is instead, as Gardner suggests, that you have a vanishingly weak reason not to pull it (2019: 43). In that case, you can change the story to one where a tree falls on the switch and activates it and ask if you think the tree's falling on the switch harmed Philip by causing Philip to smash into the second brick wall rather than the first. To me, at least, it remains just as clear that the event that switches the tracks does not harm Philip. And this intuition can't be debunked in the way that Gardner attempts to debunk the intuition in the original version of the case since trees do not act on reasons. P2 has no difficulty accounting for the fact that Victim ended up in the hospital, either, because P2 need not deny that Shooter caused Victim to be in a harmed state. Unlike the existence account of harming, P2 does not maintain that an act's causing a person to be in a harmed state is sufficient to make it the case that the act harms the person. If I push you out of the way of a bus that is about to hit you and you skin your knees as a result, P2 can agree that my act causes the harmed state that you find yourself in while maintaining that my act did not harm you since it did not make you worse off than you would have been had I not pushed you out of the way of the bus. The same would go for the case of Shooter's act. And if, as in the case of *Angry Alastair*, the police officer is going to cause something bad to happen to you either way, then P2 will say that his decision to save you \$3,000 by stealing \$5,000 from you rather than causing you \$8,000 in damages does not harm you, though, as in the case of my pushing you out of the way of the bus, his act does cause you to be in a harmed state (and, in this case, wrongfully so, since I presumably have the right to push you out of the way of the bus but he has no right to steal the \$5,000 from you; P2 entails that the police officer's choice to cause you a smaller loss rather than a larger loss does not harm you, not that his causing you the smaller loss doesn't wrong you). None of this is to say that P2 immediately produces perfectly commonsense results in all cases or that the existence account of harming does not have some advantages of its own. But on the whole, for the reasons I have tried to briefly articulate here, I am not convinced that a satisfactory solution to the non-identity problem can be found by replacing the former account of harm with the latter.

3. CHALLENGING P4: A REPLY TO KUMAR

In order to prove successful, a direct rights-based solution to the non-identity problem must provide a satisfactory answer to two questions: (1) what right is violated by the acts that take place in cases like that of Wilma or the Risky Policy? and (2) who is the bearer of the right that is violated in such cases? Put in terms of the case of Wilma, Rahul Kumar's answer to the second question is that the bearer of the right is not Pebbles herself, or potential or possible Pebbles before she was conceived, but rather "Wilma's child" in the *de dicto* sense, the sense in which Wilma's prospective child simply refers to whatever child Wilma ends up conceiving (2019: 59). And Kumar's answer to the first question is that the right that is violated in non-identity cases is not a right that a particular state of affairs not obtain, like the state of affairs in which Pebbles is blind, but rather a right that one's vital interests be recognized and given appropriate weight by others in their practical deliberations (2019: 56). Wilma, on this account, does not violate the right of Pebbles, in particular, not to be, or to be born, blind, but rather violates the right of "her child" more generally to have its vital interests recognized and given appropriate weight by Wilma when she decides when to conceive. By deciding to conceive blind Pebbles when she could instead have easily conceived sighted Rocks, that is, Wilma did not put sufficient weight on the effect that her choice would have on the vital interests of "her child" in the *de dicto* sense. This is why her act was rights violating and therefore wrong. This strikes me as an intriguing combination of answers to the two questions, one that results in a distinctive kind of rights-based solution that differs in interesting ways from those I discuss in my book. Rather than commenting on its merits relative to other rights-based approaches and to other approaches that also appeal to the *de re/de dicto* distinction (like Hare's), though, I will focus here on raising one concern about whether this approach is likely to produce a satisfactory solution to the problem.

The concern is about whether Wilma's child in the *de dicto* sense is the kind of thing that can be a bearer of rights and it can perhaps best be put in the form of a trilemma. Suppose, for example, that Wilma is a US citizen and that prior to her deciding when to conceive, I say "Wilma's child in the *de dicto* sense has a right to a US passport." One thing I might mean by this is that whatever particular child Wilma ends up conceiving, that particular child will have a right to a US passport. If she conceives Pebbles, then Pebbles will have such a right and if she conceives Rocks, then Rocks will have such a right. In this sense, there is nothing problematic about the claim that Wilma's child in the *de dicto* sense can be a bearer of rights. But in this sense, Wilma's conception decision does not seem to violate the

right that Kumar identifies. If she conceives Pebbles, then Pebbles will have a right that Wilma put appropriate weight on Pebbles' vital interests and if she conceives Rocks, then Rocks will have a right that Wilma put appropriate weight on Rocks' vital interests. But there is no reason to think that by conceiving Pebbles, Wilma failed to put appropriate weight on Pebbles' vital interests: Wilma made things as good for Pebbles as she could make them. And so there is no reason to think that Wilma's act was wrong on this first understanding of what it means to say that Wilma's child in the *de dicto* sense has rights.

Now suppose when I say that Wilma's child in the *de dicto* sense has a right to a US passport I mean something quite different. Suppose I mean that "Wilma's child" refers to a more abstract or general subject that is distinct from the particular child who fits that description, namely Pebbles, and that this more abstract or general subject-- the position that is filled by some particular child -- is itself something that has a right to a US passport. If this is what I mean, my claim seems clearly false. When "Wilma's child" is understood in this way, as the general class or category that is filled by some particular child or other, it doesn't seem like the kind of thing that could have such a right. And it's not just because it doesn't have a name and can't be photographed. It doesn't seem like the kind of thing that could have any rights. It doesn't have any of the properties that we associate with rights-bearing subjects: it isn't sentient, or rational, or autonomous, it doesn't have preferences, or desires, or interests. What would it mean to say it had a right to free speech? Or privacy, or a fair trial, or an equal consideration of its interests? But it seems that this is what it would have to mean for Wilma's child in the *de dicto* sense to have rights if Wilma's act were to be capable of violating the right to have one's vital interests given the weight they deserve. This is because Wilma's act does not fail to promote the vital interests of "her child" in the first sense, but might in some way be taken to do so in the second sense: her act fills the position "Wilma's child" with a child who has significantly less well-being than the child who would otherwise have filled the position would have had. So it seems that in the sense that "Wilma's child" understood *de dicto* can have rights, Wilma does not violate them and in the sense that Wilma might be understood as violating rights, "Wilma's child" could not have them.

But perhaps there is a third option that can help Kumar overcome this concern. Perhaps he can say that "Wilma's child" has the right in the first *de dicto* sense but that it is a right with respect to the *interests* of "Wilma's child" in the second *de dicto* sense. Annotating one summary he provides of his view (2019: 59; the parenthetical words are his while the bracketed

words are mine), we could put the suggestion like this:

“What the right entitles ‘Wilma’s child’ (understood *de dicto*) [in the first sense, as referring to whatever particular child Wilma ends up conceiving] to is that Wilma constrain her deliberations so as to make choices concerning her child [in the second sense, referring to something distinct from the particular child she ends up conceiving and more abstract and general: the position that is filled by some particular child] that are, given the information available to her, defensible as consistent with a concern for what will be intrinsically good for her child [in the second sense]”.

But this strikes me as objectionably *ad hoc*. Why would the particular child that Wilma ends up conceiving have a right with respect to how Wilma treats “her child” understood as something distinct from that particular child?

Kumar responds to a somewhat similar worry toward the end of his paper by citing the case of Bob and Sue (2019: 61), but regardless of whether his example helps to diffuse the worry he is concerned with there, it doesn’t seem to help with the worry I’m concerned with here. Bob’s right that Sue not read his e-mails, Kumar says, is not grounded in particular facts about Bob or Sue but rather in the nature of the relation in which they stand to each other, that of strangers. This might help to show how Pebbles could have some rights against Wilma that are grounded in the fact that she is Wilma’s daughter rather than in particular facts about Pebbles herself. But it doesn’t help explain how Pebbles, the specific child that Wilma ended up conceiving, could have the kind of right against Wilma that involves having a claim that Wilma treat “her child” in the second *de dicto* sense in a certain way. In order for the case of Bob and Sue to help support that claim, we would have to think not just that Bob has a right against Sue because of the stranger relationship between them, but that Bob has a right that Sue not read the e-mails of “whoever ends up fitting the description ‘stranger to Sue who has not given Sue permission to read their e-mails.’” And that seems quite implausible. If Sue reads some other stranger’s e-mails, that doesn’t violate Bob’s rights.

There is a further problem with this third option, the one on which “Wilma’s child” understood *de dicto* in the first sense has a right that Wilma give appropriate weight to the vital interests of “Wilma’s child” understood *de dicto* in the second sense when making her conception decision. The problem is that even if we can make sense of the claim that “Wilma’s child” has this right, and even if Wilma’s conceiving now rather than two months from now would violate this right, it seems reasonable for Wilma to assume that “her child” would hypothetically waive the right if she could or will

later waive it when she can. This is because if Wilma does not give significant weight to the vital interests of “her child” understood *de dicto* in the second sense, the result will not be worse for “her child” understood *de dicto* in the first sense. Whatever child she conceives today, that is, will clearly have a life worth living and would not exist at all if Wilma delayed conception by two months. Since “her child” in this sense would, or will, waive the right we are assuming it to have, the fact that “her child” has this right, if it is a fact, cannot make Wilma’s act wrong.

Kumar provides a brief response to this kind of rights waiver objection. After acknowledging that when Pebbles thinks of herself as the particular person she is, she will not resent Wilma for making the choice she made, he adds: “Which is not to say that Pebbles’ waives her right that Wilma not make the choice she did; because the right is one that constrains Wilma’s choices with respect to ‘her future child’, it is not a right that any particular future child who ends up coming into existence is in a position to waive” (2019: 60-61).

Kumar’s response here seems to depend on the claim that if a right constrains someone’s treatment of X, then only X can waive the right. Why should we accept this claim? One plausible answer to that question arises from two further claims: (1) you can only waive a right if you are the one that has the right, and (2) if a right constrains someone’s treatment of X, then only X can have that right. If only X can have a right that constrains people’s treatment of X, and if you can waive a right only if you are the one that has the right, then if a right constrains someone’s treatment of X, only X can waive that right. And if that’s so, then “Wilma’s child,” understood as whatever particular child Wilma ends up conceiving, can’t waive the right that “Wilma’s child,” understood as something other than that particular child, have its interests given due consideration by Wilma when she decides when to conceive. But Kumar cannot appeal to this argument to support the claim that his response to the rights waiver objection seems to depend on because this argument depends on the claim that if a right constrains someone’s treatment of X, then only X can have that right. And if that’s true, then “Wilma’s child” understood *de dicto* in the first sense, can’t have a right that constrains Wilma’s treatment of “Wilma’s child” understood *de dicto* in the second sense. If, on the other hand, it really is possible for “Wilma’s child” understood *de dicto* in the first sense to have a right that constrains her treatment of “Wilma’s child” understood *de dicto* in the second sense, then it is not clear why “Wilma’s child” understood *de dicto* in the first sense can’t waive that right.

While the answers Kumar suggests to the questions of whose rights are violated in non-identity cases and what rights those are may in some

respects improve on those provided by the rights-based solutions I discuss in my book, then, I'm not yet convinced that they suffice to ground a satisfactory solution to the non-identity problem.

4. CHALLENGING P4 AND P5: A REPLY TO WASSERMAN

David Wasserman's commentary on my book is largely sympathetic so it will come as little surprise that my response to his commentary is largely sympathetic, too. Rather than looking for minor points where I might be able to quibble with him, then, I will focus here on the three main arguments for the wrongness of Wilma's conduct that he discusses and add a few comments of my own.

The first kind of argument appeals to the claim that Wilma is the steward of a precious gift, the gift of life, and that by choosing to conceive Pebbles when she could instead have easily conceived Rocks, she violates a moral duty that is imposed by her stewardship of this great gift. This could be either because her decision disrespects the recipient of the gift in a way that gives the recipient a legitimate cause for complaint or because it fails to value the gift itself in a way that the gift should be valued, and so an argument of this sort might be developed as an attack on either P4 or P5 of the non-identity argument. Wasserman considers various versions of this argument and concludes that they fail to identify a wrong-making feature of an act even if the act has it, that there is no reason to think Wilma's act has the feature in the first place, or possibly both.

I have little to add to Wasserman's critique of this first kind of argument. But I would like to note one potentially important respect in which his discussion might be used to strengthen the case for the position in take in my book. Toward the end of the first part of his paper, Wasserman says that potential parents would "arguably display" an objectionable form of "indiscriminateness if they gave the gift of life to someone for whom they believed it would . . . barely be worth living. The steward of a precious gift should only give it to someone for whom it would be of significant value" (2019: 72). As Wasserman correctly notes, this kind of principle can't justify the claim that Wilma's act is wrong because the case of Wilma involves conceiving a child whose life is clearly well worth living. But, if accepted, the principle that Wasserman expresses some sympathy for here would justify the claim that Wilma's act would be wrong in a variant of the case where the doctor told her that if she conceived now her child would have a life that was just barely worth living.

Here is why this is a potentially important result. The non-identity

argument that seems to show that Wilma's act is not wrong in the original version of the story seems just as clearly to show that her act would not be wrong in this variant of the case. In biting the bullet in the original version of the story, then, I seem to be committed to biting the bullet in the variant version. And even if the strategies I use in my book succeed in making the former bullet seem not so difficult to bite (2014: 192-209), my attempt to do something similar for biting the bullet in the latter case (2014: 225-34) might not be as successful. It would therefore be helpful if, as a back-up plan, I could provide reasons to think that biting the bullet in the former case doesn't really require biting it in the latter.

I briefly mention a few such reasons in the book (2014: 226-27), but Wasserman's suggestion here provides a potentially valuable supplement. And the principle it appeals to might be supported by noting that it has significant plausibility in other cases involving the gift of life. It seems plausible, for example, to say that you are not morally obligated to donate a kidney to save the life of a stranger and plausible to say that if you do decide to donate a kidney, you may donate it to save the life of a blind person rather than the life of an otherwise comparable sighted person who would have a significantly higher overall quality of life. But it may well also seem quite plausible to say that while it isn't wrong to decline to donate a kidney at all, it would be wrong to donate a kidney to a stranger whose life is just barely worth living when you could instead at virtually no extra cost donate it to other needy people whose lives would be far, far better for them. In this respect, Wasserman's critique of the first argument against the position I take in my book also offers unexpected additional support for that position.

The second argument that Wasserman considers comes from a recent paper by Robert Noggle (2018). It's grounded in the claim that it's wrong for a person to do an act that deliberately makes it impossible for them to fulfill one of their moral obligations. Wilma has a moral obligation to protect any actual child of hers from blindness. By conceiving now rather than first taking the pill for two months, she makes it the case that she has an actual child that she cannot possibly protect from blindness. Therefore, Wilma's act is wrong. This argument, too, might be developed as an attack on either P4 or P5, depending on whether the conclusion is that Wilma's act wrongs Pebbles or that Wilma's act is simply wrong. Either way, though, as Wasserman correctly points out, the argument is unsuccessful because Wilma's act does not make it impossible for her to fulfill an obligation she would otherwise have had. Prior to deciding when to conceive, it was already impossible for Wilma to protect Pebbles from blindness. Wasserman goes on to note a few other problems with this argument, but I think it's worth adding a few more that he doesn't discuss.

First, if Noggle's argument shows that Wilma's act is wrong, then it also shows that it's wrong for Betty to conceive Bamm-Bamm when she knows that any child she conceives will be blind. In whatever sense Noggle could claim that Wilma is responsible for making it impossible for her to protect Pebbles from blindness, Betty is just as responsible for making it impossible for her to protect Bamm-Bamm from blindness. But, as I noted in the context of Gardner's solution, the implication that it's wrong for Betty to conceive a blind child when she knows that any child she conceives will be blind renders a proposed solution to the non-identity problem unable to satisfy the modesty requirement.

Second, and perhaps more importantly, a parent's obligation to protect their actual children from blindness can be overridden. And one scenario in which the obligation is clearly overridden is that in which a child's losing their vision is necessary for them to continue living. If a child will die without life-saving surgery that would cost them their vision, for example, it would clearly not be wrong for the parents to authorize the surgery. And one very plausible way to account for this is to say that if the child understood the situation, they would consent to the act that would cause their lack of vision or that later, when they do understand the situation, they will give their subsequent consent to that act. But if we can say this about the surgery case, it's not clear why we can't also say it about the case of Wilma: even if Noggle were right to think that Wilma's act makes it impossible for her to fulfill her obligation to protect Pebbles from blindness, this would still not be enough to make the act wrong because Pebbles would, or will, consent to relieve Wilma of this obligation¹.

The third kind of argument that Wasserman considers maintains that Wilma's act is wrong because it involves insensitivity to the harm of blindness. This is the most developed and detailed part of Wasserman's paper, but his considered position on this kind of argument is somewhat less clear. In the introduction, he says of the three kinds of argument he considers "I doubt that Wilma acts wrongly in either of the first two ways and remain uncertain about the third" (2019: 67). In the section of the paper he devotes to the argument, he considers a variety of ways that someone might defend some version of the argument and responds in ways that I, at least, took to be good reasons not to accept any of them. But in the conclusion, when he looks back on that section, he says in part that "I found it plausible to claim that [Wilma] displayed insensitivity to the

¹ At the risk of engaging in just one minor quibble with Wasserman: for the same reason, I think Wasserman is mistaken, in the context of his discussion of the first argument, to suggest that the 14-year-old girl's act of conceiving is wrong in Parfit's case because it imposes an unacceptable risk that she would later be unable to fulfill her obligations to the child (Wasserman 2019: 70).

(stipulated) hardship of blindness in choosing to have a blind child without seriously considering what life would be like for such a child” (2019: 82). And he adds that “I also found that such insensitivity would give Pebbles, the blind child she ended up having, grounds for complaint” (2019: 82), which means this third kind of argument would best be understood as aiming at P4. While Wasserman at times seems close to neutral about this kind of argument, at times thoughtfully skeptical, and at other times more sympathetic, I will simply conclude with a few thoughts about this kind of argument to the extent that he seems favorably inclined toward it.

First, I don’t think the fact that Wilma makes her choice based on convenience means she is not sufficiently sensitive to the harm of blindness. Suppose you can save one but only one of two drowning children, that it is a bit more convenient for you to save the one who is a bit closer to you, and that the two children are otherwise comparable except for the fact that the one who is a bit closer to you is blind. You can be fully sensitive to the respects in which the blind child’s blindness has a negative impact on his well-being while still recognizing that his life is well worth living and while thinking that it isn’t morally wrong to let convenience be the tie breaker in determining which life you will save. There may be other differences between life creation and life extension cases, but I don’t see why deciding based on convenience in the former case would have to involve this kind of insensitivity if deciding based on convenience in the latter case would not.

Second, suppose that Wilma’s act really was based on her being insensitive to the harm of blindness. Suppose, for the sake of the argument, that she didn’t give a moment’s thought to the kind of life her blind child would end up leading. Even if this is so, this would not show that her act was wrong, only that her motive was objectionable. Suppose, for example, that a man agreed to donate a kidney to his young daughter, who would have died without it, only because he knew his wife would be angry at him if he refused. Clearly his child would later have good reason to complain about his lack of concern for her well-being. But, just as clearly, this does not mean that it was wrong for him to donate the kidney to her.

Finally, even if we thought insensitivity of this sort could suffice to make the act itself wrong, and not just make the motive for doing the act objectionable, this would still not be enough to solve the non-identity problem because, if necessary, we could simply stipulate that Wilma does not act with this kind of insensitivity. Perhaps she had read my book carefully and was convinced by it, and so decided to conceive now only because she could see that despite whatever hardships would ensue for Pebbles, Pebbles would not be harmed or wronged by her doing so. In the

end, then, I don't think any of the arguments Wasserman discusses provide a promising route to solving the non-identity problem by attacking either P4 or P5, and I am, if anything, more skeptical of some of them than he is.

5. CHALLENGING P5: A REPLY TO MULGAN

Tim Mulgan's challenging paper can be summarized in terms of three claims: (1) that I can't account for the wrongness of failing to prevent extinction in his Starship case (2019: 86) without also conceding the wrongness of Wilma's act, (2) that Rule Consequentialism (RC) can account for the wrongness of failing to prevent extinction in his Starship case, and (3) that the way that RC does this enables it to show that Wilma's act is wrong without committing itself to any unacceptably immodest further conclusions. The second claim clearly seems right, but I'm not convinced of the first or the third.

Let's start with the first. There is at least one argument for the wrongness of failing to prevent extinction in the Starship case that strikes me as plausible and that does not entail that Wilma's act is wrong. The argument is grounded in two claims. The first is that if you can prevent the destruction of something of great value at a relatively little cost, it is wrong for you not to do so. The second is that the human species has great value. If both these claims are true, then it follows straightforwardly that it would be wrong not to prevent the extinction of the human species in the Starship case. And since choosing to conceive Pebbles rather than Rocks does not involve Wilma failing to prevent the destruction of something of great value – indeed, by conceiving Pebbles rather than Rocks, she does not fail to prevent the destruction of anything at all – this argument for the wrongness of failing to prevent extinction in the Starship case does not entail that Wilma does something wrong by conceiving Pebbles rather than Rocks.

Mulgan might object that this argument violates my independence requirement. “[J]ust as Boonin rejects ad hoc solutions tailored to avoid the Implausible Conclusion [that Wilma's act is not wrong],” he writes, “we must also reject ad hoc solutions to the Nihilistic Conclusion [that failing to prevent extinction in the Starship case is not wrong]” (2019: 89). But this argument does not violate that requirement. Put in terms of the case of Wilma, the independence requirement says that we can't justify accepting a claim simply because accepting it produces the result that Wilma's act is wrong. When applied to the case of Starship, then, the requirement says that we can't justify accepting a claim simply because accepting it produces the result that failing to prevent extinction in the Starship case is wrong. But each of the two claims the argument depends on seem reasonable

independent of the fact that, when combined, they produce this result. The first seems to best explain the judgments we would make in a wide range of cases where something of great value can be easily saved from destruction, and the second seems to best explain the judgments we would make in a wide range of cases where we are asked to judge the value of the human species against the value of something else.

It's true, of course, that the first claim this argument appeals to is incompatible with P5 of the non-identity argument, since P5 says that if an act wrongs no one then the act is not wrong. But in order to solve the non-identity problem, it's not enough to show that one of the premises of the non-identity argument is false. One must show that the reason it's false can be applied to cases like that of Wilma. And if P5 is false because an act can be wrong without wronging anyone when it fails to easily prevent the destruction of something of great value, this won't help to show that Wilma's act is wrong. Again, Wilma's act does not involve the failure to prevent the destruction of anything, let alone the destruction of something of great value. So I don't think the claim that failing to prevent extinction in the Starship case would be wrong poses a problem for the claim that Wilma's act is not wrong.

What about Rule Consequentialism? RC clearly has no problem accounting for the wrongness of failing to prevent extinction in the Starship case. But I'm not convinced by Mulgan's claim that RC can show that Wilma's act is wrong without committing itself to any unacceptably immodest further conclusions. Either RC permits Wilma to conceive Pebbles or, if it doesn't, it forbids too much to satisfy the modesty requirement. If we set aside doomsday scenarios like the Starship case, Mulgan seems to concede that there's no particular reason to think RC would deem Wilma's act to be wrong. An ideal code capable of being internalized by people can't be too demanding or impersonal, and a code that isn't too demanding or impersonal will permit a variety of sub-optimal choices. Wilma's would plausibly be among them. But, he argues, once the threat of extinction is added to the equation, "RC must abandon its commitment to a liberal ideal code that incorporates a permissive procreative ethic and a presumption of reproductive freedom" (2019: 95). The result, instead, is that "RC will now select a much less permissive ideal code that rejects any presumption of reproductive freedom and also includes a much stronger obligation to do whatever one can to enhance future well-being. When making a procreative choice, the RC ideal agent will not feel free to consult her own desires when this would lead to a sub-optimal result in terms of future well-being" (2019: 96). Clearly such an agent would not feel free to make the choice that Wilma makes.

Mulgan considers the possibility that the ideal code would be more fine-grained than this, leaving people free to make the choice Wilma makes in a world like ours while restricting their freedom in much bleaker scenarios. But he argues that internalizing these more flexible versions of the code across entire generations would produce suboptimal results for three reasons: making the code more complex makes it harder to internalize, we might not be good enough at recognizing when an extinction threat is upon us and thus disastrously fail to shift behavior toward the more restrictive rules when we were supposed to, and the extra benefit in terms of increased freedom for us is not enough to outweigh these potential costs to future generations (2019: 96). As with many attempts to apply consequentialist reasoning to complex concrete questions, it's hard to know just what to say about the claim that the ideal code would have to be as rigid and illiberal as Mulgan claims it would be in this respect. But let's suppose that he's right.

If he is right, there is no reason to limit the code's demandingness to procreative decisions. In the face of a genuine threat of extinction, virtually every kind of decision a person makes could have potentially disastrous consequences for the survival of the species: decisions about what and how much to consume, what and how much to produce, what and how much to learn, who to interact with, how much time to set aside for leisure, and so on. If the factors that Mulgan appeals to really suffice to show that an ideal code would not allow Wilma to "feel free to consult her own desires when this would lead to a sub-optimal result in terms of future well-being" when making procreative decisions, then, they will also suffice to show that an ideal code would not leave her free to consult her own desires when this would lead to a sub-optimal result in terms of future well-being when making consumption decisions, production decisions, educational decisions, socializing decisions, leisure time decisions, and so on. Indeed, it becomes difficult to see how the ideal code could leave people free to consult their own desires when making any decision at all if their choice would lead to a sub-optimal result in terms of future well-being. For all intents and purposes, people would not feel free to do anything just because they wanted to do it, if doing it would produce anything but optimal results for future well-being. The "moderate, liberal alternative" that RC was supposed to offer in place of act-consequentialism seems to have been left far behind.

This result seems worrisome enough even if it's limited to cases like Wilma's, where people have a relatively modest desire to act sub-optimally. What about cases where they would be more significantly burdened by not feeling free to act sub-optimally? Mulgan treats my case of Jane as an

example of this and suggests that RC could appeal to the difference in burdens in the two cases to explain why Jane should feel free to act sub-optimally while Wilma should not. Because of the more significant burden to Jane, he suggests, “Only in extreme circumstances will the RC ideal agent feel obliged to have a child she doesn’t want. As Jane’s circumstances are not extreme, she does no wrong” (2019: 99). But this suggestion is difficult to square with Mulgan’s earlier claim that because of the risk of extinction, RC’s ideal code “rejects any presumption of reproductive freedom.” If there is no presumption of reproductive freedom, why should the cost to Jane in reproductive freedom count against requiring her to make the optimal procreative choice? And, perhaps more importantly, if Jane can be relied on to recognize whether she is in extreme circumstances, why can’t Wilma? If Jane can successfully internalize a code that permits her to act sub-optimally in non-extreme circumstances while requiring her to act optimally in extreme circumstances, why can’t Wilma? If Mulgan’s argument succeeds in showing that the ideal code would leave Wilma feeling unfree to make the sub-optimal procreative choice that she makes, then, not only does it also show that we should virtually never feel free to act sub-optimally just because we want to, but it will require us to make optimal choices across a broad range of domains even in cases where we have more than just a modest desire not to do so. The implications of RC have become extremely illiberal by this point.

Let’s suppose I’m right about this. Is that a problem? It depends on what we’re trying to do. If we are trying to arrive at the correct moral theory, it need not be a problem. Perhaps the correct moral theory has a number of implications that will strike most people as extremely implausible. If there is a good enough reason to think that RC is true and a good enough reason to think that RC’s ideal code should leave Wilma feeling unfree to conceive Pebbles rather than Rocks, then there is a good reason to think that the correct moral theory tells us to virtually never feel free to act sub-optimally just because we want to or even if we have a significant interest in doing so. I think that implication will strike most people as highly implausible, but if there’s a good enough reason to accept it, then we should accept it despite its being so implausible.

But if what we are trying to do is solve the non-identity problem, then I think the extremely illiberal implications of the version of RC that we would have to accept in order to show that Wilma’s act is wrong really are a problem. The non-identity problem arises in the first place, after all, only because the conclusion of the non-identity argument strikes so many people as implausible. The claim that Wilma’s act isn’t wrong is not self-contradictory or empirically false; it just seems implausible. The problem

that stands in need of resolution is that of avoiding being committed to an implausible claim. If the means of avoiding the implausible claim commit us to an even more implausible claim, then, this does not really solve the problem. And if the only way to show that Wilma's act is wrong really does commit us to the claim that we should almost never feel free to act sub-optimally in terms of future well-being just because we want to or even if we have a significant interest in doing so, it seems to me that the means of avoiding the claim that Wilma's act is not wrong commits us to an even more implausible claim. Not only is there a plausible explanation of the wrongness of preventing extinction in the Starship case that does not commit us to saying that Wilma's act is wrong, then, but the extremely illiberal version of RC that seems required to say the acts are wrong in both cases seems considerably less plausible than simply allowing that Wilma's act is not morally wrong.

6. FROM TWO-OPTION CASES TO THREE-OPTION CASES: A REPLY TO ROBERTS

I have focused to this point on the case of Wilma. What about the case of Risky Policy? I think the non-identity problem as it arises in the former case is fundamentally the same as the problem as it arises in the latter case. If my response to the problem works well in the one, then it works well in the other. But Melinda Roberts thinks there's an important difference between the two. Wilma, she argues, is a genuine two-option case, but Risky Policy is really a three-option case. And this means, according to Roberts, that even if my treatment of Wilma proves satisfactory, it can't be successfully extended to the case of Risky Policy. Extending my analysis of the former case to the latter, she argues, generates problems for my defense of P2 of the non-identity argument and for my strategy of trying to make it seem sufficiently plausible that we ultimately accept the conclusion of that argument. The first concern has to do with a claim about harm while the second concern has to do with a claim about wrongness. I will respond to these two concerns in turn. In doing so, I will set aside the question of whether Roberts is right to think the Risky Policy case really should be understood as a three-option case and focus my response on two of the three-option cases she uses in developing her arguments: Medical Case I and the case of Warren and Percy.

Let's begin with a three-option case that Roberts thinks poses a problem for P2, the case she calls Medical Case I (2019: 114). In this case, you are dying, and your doctor has three options: (1) do nothing, in which case you will die, (2) give you medicine A, in which case you will survive and partly

recover with a life clearly worth living but at a significantly reduced level of well-being, or (3) give you medicine B, in which case you will survive and fully recover and so enjoy a significantly higher level of well-being than you'll enjoy if the doctor gives you medicine A. In the case as Roberts describes it, the doctor gives you medicine A. But it turns out that if he had not given you medicine A, he would not have treated you at all.

P2 endorses the counterfactual account of harm: if A's act harms B, then A's act makes B worse off than B would otherwise have been. Roberts thinks Medical Case I poses a problem for this account of harm because it entails that the doctor's act does not harm you and because she thinks this result is implausible (2019: 123). P2 entails that the doctor's act does not harm you because his act does not make you worse off than you would otherwise have been. Because of his act, you survive and partly recover with a life clearly worth living. If he had not done this act, he would instead have done nothing, in which case you would have died. Surviving and making a partial recovery is not worse for you than dying. And so, according to P2, the doctor's act of giving you medicine A does not harm you. The question is whether this implication poses a problem for P2.

I don't think it does for two reasons. First, if we preserve the structure of the case but remove it from the medical context, I doubt people will think the act harms you. Suppose I'm considering doing you no favor, a small favor, or a big favor. I decide to do you a small favor and had I decided not to do you the small favor, I would have done you no favor. Do I harm you by doing you a small favor in this case? I doubt people will think I do. So if P2 entails that I don't harm you in this case, that shouldn't be a problem for P2. The medical case is different from this case only because the doctor is obligated to give you the best available treatment while I am not obligated to do you the biggest available favor. So we can explain away our initial belief that the doctor harms you when he gives you medicine A, assuming we had that belief in the first place, by saying that what we were really responding to is the fact that the doctor wrongs you by making you worse off than he is obligated to make you. And saying that is consistent with accepting P2. His act of giving you medicine A doesn't harm you. It benefits you. But because it doesn't benefit you as much as he is obligated to benefit you, he wrongs you nonetheless by doing it.

But suppose you think it really is a problem for P2 if it can't account for the claim that the doctor harms you by giving you medicine A. In that case, we can vindicate P2 by explaining how it can account for this claim. We can do this by breaking the three-option case down into a pair of two-option cases. Doing so will also help to show why I don't think three-option cases introduce any distinctive problems that can't be successfully

addressed in terms of my treatment of two-option cases.

Here's what I have in mind. First, the doctor chooses between giving you medicine B and not giving you medicine B. From these two options, he chooses not to give you medicine B. That decision makes you worse off than you would otherwise have been. Had he not made this decision, the best outcome would still be available to you. Because he made this decision, the best option is no longer available to you. So P2 says that he harms you by taking that option off the table. And that implication seems correct. Having decided that he's not going to give you medicine B, he now makes a second decision, choosing between giving you medicine A and doing nothing. He decides to give you medicine A. Because of his prior decision, it's his fault that if he hadn't given you medicine A, he would have done nothing. And because it's true that if he hadn't given you medicine A he would have done nothing, it's also true that this second choice – the choice to give you medicine A rather than to give you nothing – does not itself make you worse off than you would otherwise have been. So P2 entails that he does not harm you by giving you medicine A rather than doing nothing. But this implication, too, seems correct. The doctor harms you by taking option (3) off the table but does not harm you by choosing option (2) over option (1) given that option (3) is off the table. That's what P2 says and that seems right. So I don't think three-option cases pose a problem for the account of harm I endorse in the book.

What about my account of wrongness? Roberts presses this problem in the context of the case of Warren and Percy (2019: 107). A sperm and an egg are in a glass dish, moments away from fertilization. Warren has three options: (1) throw away the contents of the dish, in which case Percy will never exist, (2) put a chemical in the dish and leave it there, in which case Percy will exist with some significant impairment but with a life clearly worth living, (3) leave the dish alone, in which case Percy will exist without the impairment and have a considerably better life. Warren chooses the second option. If he hadn't, he would have chosen the first. Roberts argues that the approach to the non-identity problem that leads me to accept the conclusion that Wilma's act is not wrong also commits me to accepting the conclusion that what Warren does isn't wrong. And while, unlike many philosophers, Roberts is sympathetic to my accepting the claim about Wilma, she takes it as clear that we shouldn't accept the parallel claim about Warren. I think that's right, but I also think there's nothing about my treatment of the Wilma case that prevents me from saying so. And so, with respect to wrongfulness as with respect to harmfulness, I don't think three-option cases introduce any problems that can't be addressed in terms of the way I address two-option cases like that of Wilma.

We can again see this by analyzing the three-option case in terms of a pair of two-option cases and then looking at the question of whether Warren's behavior is wrong. First, Warren chooses between leaving the dish alone and not leaving the dish alone. He decides not to leave the dish alone. Next, having decided that he's going to tamper with the dish, he chooses between two ways of tampering with it: tossing out its contents, or pouring the chemical into it. From these two options, he chooses to pour the chemical into the dish. If he hadn't poured the chemical into the dish, he would have thrown away the contents of the dish. As a result of Warren's choices, Percy comes to exist with the significant impairment.

What does my treatment of the two-option case of Wilma commit me to saying about the three-option case of Warren? Consider first his decision to tamper with the dish rather than to leave it alone. Assuming that tampering with the dish doesn't wrong anyone else, such as the sperm or egg donor or those providing the technical support, my treatment of the case of Wilma doesn't yield a determinate answer to the question of whether this choice is wrong. If choosing to tamper with the dish leads him to toss out the contents, then Percy will never exist and so no one will have been wronged or harmed by his decision. If, on the other hand, his decision to tamper with the dish leads him to pour the chemical into it, then it will turn out that his decision to tamper with the dish did harm and wrong someone by making Percy worse off than he would otherwise have been. This implication is a bit complicated, but not in a way that should make it seem problematic. You can't decide whether someone acts wrongly by planting a bomb in an abandoned building they own, for example, without knowing whether they subsequently lure someone into the building before the bomb explodes. So I don't think my analysis generates any problems in assessing the wrongness of Warren's first choice. That choice is wrong if it leads him to subsequently put the chemical in the dish and not wrong if it doesn't.

Now let's focus on the second choice. It is now the case that Warren will either pour the chemical into the dish or throw away the contents of the dish. He chooses to pour the chemical into the dish rather than throw away the contents of the dish. My analysis commits me to saying that he does not do anything wrong by making this second choice. That result also seems right to me. And as long as we keep in mind that he has already taken the option of leaving the dish alone off the table, I think it will seem right to you, too. Suppose, for example, that the only way he could prevent the contents of the dish from falling into a garbage can was by pouring the chemical into the dish. Choosing to pour the chemical into the dish in that case would not seem wrong. If it's his own decision to tamper with the dish rather than contingent circumstances that limit him to these two options,

that doesn't seem to change that fact that choosing to pour the chemical rather than to throw away the contents doesn't by itself wrong or harm anyone. So is Warren's behavior in this case wrong? Yes. His decision to tamper with the dish ends up harming Percy. If his decision to tamper with the dish had instead led him to throw out its contents, then it would not have been wrong. These results seem right to me and are entailed by breaking the case down into a pair of two-option cases and then treating the two-option cases in the way that I treat the two-option case of Wilma in my book. It still seems to me, then, that if my treatment of the two-option case of Wilma in the book is satisfactory, it can be successfully extended to the case of Risky Policy even if Roberts is correct that that case must be understood as a three-option case.

BIBLIOGRAPHY

- Boonin, D., 2014: *The Non-Identity Problem and the Ethics of Future People*, Oxford: Oxford University Press.
- Gardner, M., 2015: "A Harm-Based Solution to the Non-Identity Problem. *Ergo* 2: 427-44.
- 2019: "David Boonin and the Non-Identity Argument: Rejecting the Second Premise", *Law, Ethics and Philosophy* 7: 29-47.
- Hare, C., 2007: "Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?", *Ethics* 117: 498-523.
- Kumar, R., 2019: "Rights, Wronging, and the Snares of Non-Identity", *Law, Ethics and Philosophy* 7: 48-62.
- Malek, J., 2019: "The Non-Identity Non-Problem: Constructing Continuity of Identity among Possible Persons", *Law, Ethics and Philosophy* 7: 15-28.
- Mulgan, T., 2019: "How Rule Consequentialism Avoids Boonin's Implausible Conclusion", *Law, Ethics and Philosophy* 7: 84-101.
- Noggle, Robert. 2018. "Impossible Obligations and the Non-Identity Problem", *Philosophical Studies*.
- Roberts, M., 2019: "Boonin on the Nonidentity Problem: Too Narrow on the Facts? Too Narrow on the Theory?", *Law, Ethics and Philosophy* 7: 102-126.
- Wasserman, D., 2019: "Chasing the Elusive Wrongdoing Intuition", *Law, Ethics and Philosophy* 7: 63-83.