

Why Wear Blinders? Boonin and the Narrow Approach to the Non-identity Problem*

MELINDA A. ROBERTS

The College of New Jersey

ABSTRACT

Boonin endorses reasoning that leads to what he calls the Implausible Conclusion regarding when future-directed choices that at first glance seem to impose substantial burdens on future people are permissible. According to the Implausible Conclusion, such choices – provided they serve as “but for” causes of the future person’s coming into existence, and provided they don’t harm still other people – turn out to be permissible. That’s so, according to Boonin, since those same choices don’t in fact impose the substantial burdens on future people that we at first glance think they do and on the assumption (an assumption not questioned here) that any plausible moral theory will be person-affecting in nature and thus limit wrongdoing to choices that harm existing or future people. The purpose of this paper is to argue that the reasoning Boonin urges us to accept fails to establish the Implausible Conclusion, on two closely related grounds. First, Boonin ignores some of the morally significant factual details of many of the cases he wants his Implausible Conclusion to hold for, including some of their critically important *modal* details. And, second, at least one principle that Boonin endorses – a principle that governs when a future person is harmed by a prior choice – is clearly false, even on the assumption (also not questioned here) that harming a person is a matter of making that person worse off. The combination of those two mistakes results in an inappropriately narrow approach to the non-identity problem. Jettisoning that approach means that we are free, thankfully, to say what we thought we should say about at least many versions of the non-identity problem to begin with: that the Implausible Conclusion is false, and that future-directed choices that seem at first glance to impose substantial burdens on future people are indeed often wrong.

* I am very grateful to the anonymous referees for their invaluable comments on a prior version of this paper.

Keywords: Risky policy case, counterfactual account of harm, probability and harm, medical treatment cases and harm, contrast between two- and three-option cases.

1. WHAT IS THE NON-IDENTITY PROBLEM?

1.1 *The Implausible Conclusion.* David Boonin's presentation of the non-identity problem is sufficiently clean and clear that any interested reader – philosopher or not – will easily grasp both the reasoning Boonin identifies as giving rise to the problem and how difficult it is to say just where that reasoning goes awry. Putting things that way, of course, may seem to suggest that reasoning does go awry. This is certainly the conventional view. But it's exactly this conventional view that Boonin contests. The main argument of his book is that that reasoning *doesn't* go awry.¹ According to Boonin, to solve the problem is to accept the reasoning. Specifically, it is to accept that the non-identity problem isn't a *problem* after all but rather a sound argument to a wide range of surprising results. Boonin collects those results under the heading the *Implausible Conclusion*, and then ices his cake with some reasons to think that those results are not really "so implausible after all."²

Thus we thought that it was a problem for a moral theory that it evaluates as morally permissible all sorts of choices that impose devastating burdens on future people. But that was a mistake. The better view, according to Boonin, is that choices that burden the future people they bring into existence – *existence-inducing choices* – are perfectly permissible provided that those future people's lives are on balance worth living and that still other people – whether existing or future – are not themselves also burdened by those choices.³

That conclusion is sweepingly exculpatory. That's so in virtue of the fact that bringing a new person into existence – causing a new person to exist – is not as hard to do as we might have thought. It's not necessarily a

1 Boonin (2014), chaps. 1 and 7.

2 Boonin (2014): 191.

3 "Existence-inducing choices," I should note, is my term, not Boonin's. I should note that another point taken for granted in discussions of the non-identity problem is that the proper evaluation of a given choice must take into account not just the burdens that come with that choice but also the benefits. Thus, to properly evaluate whether the choice to vaccinate a child is morally permissible, we must balance the burdens against the benefits. The pain the child experiences is a burden; that the child is then protected against, for example, measles is a benefit. By hypothesis in the non-identity cases, the benefits always do at least counterbalance the burdens and the life is, despite the burdens, on balance always clearly worth living.

matter of getting pregnant or impregnating someone else. Rather, as Derek Parfit and others have argued, choices that impose burdens on future people and seem to have nothing to do with reproduction, on closer inspection, often – indeed, *typically* – turn out to existence-inducing.⁴

How can that be? How, for example, can a choice that we make today stripping distant future generations of important natural resources also serve to bring just those same people into existence?⁵ Here, we take for granted that any distinction in the gametes involved in the conception of any one person inevitably would lead to a distinction in the identity of the person conceived. Then, as Parfit has argued, even the slightest departure from a given sequence of acts and events – any little change that might affect the “timing and manner” of conception, even changes aimed solely at avoiding or mitigating a given burden on behalf of a given future person – will lead, or at least, as Gregory Kavka observed, *very probably* will lead, to that person’s never coming into existence at all.⁶ How many of us would have existed, as Parfit puts it, had “railways and motor cars” never been invented?⁷ Had World War II not been fought? Had our parents – indeed our forebears; indeed the prehistorical sea-creatures from which we have ultimately descended – not had sex just when and how and with whom they in fact did?

Now, an *impersonal* approach to the evaluation of choice may deem the fact that a given choice brings a particular person into existence beside the point. That a choice that has burdened a person has also conferred on that same person a counterbalancing benefit – the gift of life itself – is not morally relevant. If an *alternate* choice would have avoided the burden by way of bringing a *distinct* person into existence – a person *non-identical* to the person burdened by the choice in fact made – then other things equal that is reason enough, under an impersonal approach, to deem the choice in fact made morally wrong.

But suppose that we find a *person-affecting*, or *person-based*, approach to the evaluation of choice more intuitive than an impersonal approach. Specifically, suppose we accept what is called the *person-affecting*, or *person-based*, *intuition*, according to which we do nothing wrong, other things equal, in failing to bring a person into existence to begin with. Suppose we accept that what is “bad” must be “bad for” someone – that is, that a choice can be morally wrong *only if* that choice makes things worse for – or *harms* – at least some existing or future person.

4 Parfit (1987), Part IV; Kavka (1982).

5 This is Parfit’s Depletion case. See Parfit (1987): 361-363.

6 Kavka (1982): 100 n. 15.

7 Parfit (1987): 361.

It's then that we face the non-identity problem. It's then that we seem to find ourselves compelled to say that the existence-inducing choice under scrutiny – however burdensome, and provided just that the life is worth living and no one else is affected – is perfectly permissible.

Thus does Boonin endorse a line of reasoning that ends in a sweepingly exculpatory position in respect of the choices that we make today that will predictably burden – and perhaps severely burden – the future people they bring into existence. According to Boonin, we are, in other words, compelled to say that when it comes to future people, almost *anything goes*. That's the *Implausible Conclusion*. And that's the line of reasoning and the conclusion that Boonin thinks we must accept in order to solve the non-identity problem.⁸

1.2 *The Wilma and Pebbles case*. Boonin explores both *direct* versions of the non-identity problem, including the Wilma and Pebbles case, and *indirect* versions of the problem, including the Risky Policy case.⁹ A *direct* version of the problem is one in which a “choice directly determines which particular person will exist after the choice is made.”¹⁰ In the indirect version of the problem, a “choice has consequences that initiate a complex chain of events that eventually have [a] decisive effect on which particular people exist”¹¹

However, as Boonin understands things the underlying logic of the problem is the same in the Wilma and Pebbles case and in the Risky Policy case. Now, my own view is that we should closely question whether that's in fact so. It's not that there's a clear moral distinction to be made between Boonin's *direct* cases and his *indirect* cases but rather that not all non-identity cases are themselves of a piece. If that's so then – contrary to Boonin's view – the correct account of the Wilma and Pebbles case and the

8 As noted earlier, the icing on his case are the reasons he gives, toward the end of his book, that suggest that the Implausible Conclusion is not really so implausible after all. I won't try to evaluate those reasons here, other than just to note that he proceeds in part by analogy. According to Boonin, an agent is “within his rights” and does not “do anything positively immoral” when he or she fails to rescue a child from certain death by starvation in a case where, at little cost to the agent and at no cost to anyone else, the agent easily could have rescued the child. If we let the agent off the hook in this case – the famine relief case; Boonin (2014): 192-195 – why is it so hard to believe we should let the agent off the hook in the non-identity cases as well? The problem with Boonin's argument is that it's not at all clear that we do let the agent off the hook in the failure to rescue cases. The transitivity argument he presents later in the same chapter is no more convincing. Boonin (2014): 198-205.

9 Boonin (2014): 2-7. As Boonin acknowledges, his Wilma and Pebbles case parallels Parfit's two medical programs case. See Parfit (1987): 367-368. Boonin's Risky Policy case is a version of Parfit's risky policy case. Parfit (1987): 371-372.

10 Boonin (2014): 5.

11 Boonin (2014): 5.

correct account of Risky Policy may be quite distinct. I will return to that point in what follows.¹² For introductory purposes, however, to understand the non-identity problem *as Boonin himself understands the problem*, it's enough to consider just the Wilma and Pebbles case.

There, the child Pebbles is born “incurably blind” as a result of a condition that afflicts her mother, Wilma.¹³ Pebbles' life is clearly worth living. Wilma could have taken pills for two months that would have cured the condition. She could have then conceived a child – Boonin calls him Rocks – who would have been fully sighted. The catch is that Pebbles and Rocks are distinct (non-identical) individuals. In no clear sense would it have benefited *Pebbles* for Wilma to have delayed conception and produced *Rocks* instead. By hypothesis, no one else will be affected by how Wilma makes her choice. If – citing the person-based intuition – we think that wrongdoing is a matter of burdening, or making things worse for, some existing or future person, we seem to find ourselves compelled to say that what Wilma has done, in not taking the pills and producing a blind child when she could have taken the pills and (later) produced a sighted child instead, is perfectly permissible.¹⁴

Is this a non-identity *problem*, a line of reasoning that generates a result we must somehow get around? Or is it instead a sound non-identity *argument*, an argument that shows that Wilma's choice is perfectly permissible?

Boonin takes the latter view. And the evaluation of permissibility – to my ears, in any case – in his Wilma and Pebbles case does not seem implausible. At least, to my ears, the example does not on its own establish that the person-based intuition is false. When a given choice *maximizes* wellbeing for all involved – for all, that is, the existing and future people who exist under that choice – is it really that clear that it's *morally wrong* to bring the imperfect child over the more perfect child into existence? I'm not sure that it is – and, if it's not, then we don't have a clear-cut counterexample establishing the person-based intuition itself as false.

But neither, of course, does any one case prove that any given principle is true. The difficulty for Boonin is that still other cases suggest that his approach to the non-identity problem – his focus on certain factual details of the various non-identity cases and disregard of others, in combination with his formulation of the person-based intuition itself – really does at some point go awry.

12 See parts 3 and 4 below.

13 Boonin (2014): 2-3.

14 Boonin (2014): 5.

Thus consider the Warren and Percy case. Warren wants to have a child, and he wants to use the technique of in vitro fertilization to achieve that end. With technological support provided by others, Warren has arranged things so that his own single sperm and a single donor egg are in close proximity to each other in a glass dish and just moments away from fertilization. During that interval, Warren introduces an iatrogenic chemical into the dish, a chemical that Warren understands will inevitably cause the resulting child – call him Percy – to suffer a certain impairment. Why? Perhaps Warren is eager to test “accepted wisdom” which has it that the chemical is harmful; perhaps Warren has agreed earlier on to produce a child for a couple who seeks to raise a child with exactly the impairment the chemical will cause; perhaps it’s just that Warren has been paid \$100 to introduce the chemical into the dish. Whatever the details, it’s a stipulation of the case that, had Warren not introduced the chemical into the dish, Warren would have simply discarded the contents of the dish and Percy would never have existed at all. Fertilization proceeds, the embryo is then transferred to the uterus of a commercial surrogate, the child Percy eventually develops out of that embryo and is born. Percy’s existence is predictably flawed but nonetheless clearly worth having. And it’s an existence that Percy never would have had to begin with had Warren not chosen just as he did.

It’s a difficulty for Boonin that his approach to the non-identity problem – that is, the very approach that generated the result that *Wilma’s* choice is permissible – now seems to generate the result that *Warren’s* choice is permissible. In contrast to the result that *Wilma’s* choice is permissible, the idea that *Warren’s* choice is permissible seems *irredeemably* implausible. It seems then that after all Boonin’s approach to the non-identity problem – at either the factual level or the theoretical level or both – has gone awry.

1.3 *Plan for this paper.* We haven’t yet laid out the principles that Boonin himself adopts for purposes of formulating the person-based intuition, the idea, that is, that what is “bad” must be “bad for” someone. That work, which we shall turn to in part 2 below, will help to show why Boonin thinks we are committed to the Implausible Conclusion.

I suggested above that we do, after all, have reason to think that Boonin’s approach to the non-identity problem at some point goes awry. The main purpose of this paper, then, will be to explore just how that might have happened.

Part 3 below examines one possible explanation – that Boonin himself hasn’t fully appreciated the moral significance of some of the factual details of his own cases. Now, that’s not a mistake he seems to make in the *Wilma* and *Pebbles* case. But in other cases, including, as we shall see,

Risky Policy, certain factual details – and specifically *modal* details regarding what the agent's *available alternative* options are just prior to choice – are never mentioned and are certainly never explored. Yet – as I will argue – it is a mistake to put blinders on in respect of those details. It is a mistake to assume in advance that those details have *no moral relevance* at all to our evaluation of the choice under scrutiny.

Perhaps, however, Boonin's difficulty arises not out of a failure to appreciate fully the moral significance of the factual details of the various cases but rather out of the particular principles he sets forth for the purpose of articulating the person-based intuition. Part 4 below explores that alternate explanation for the difficulty Boonin finds himself in. There, I will argue that principles that are more sensitive to the factual details of the various non-identity cases – principles that recognize the moral significance of those details rather than simply setting them aside altogether – may enable us to make the distinctions we need to make among the various cases.

The two possible explanations are clearly linked. My argument thus will be that Boonin's understanding of the relevance of the facts of his own cases in combination with the principles he adopts for purposes of understanding the person-based intuition and applying that intuition to those cases reflect an *inappropriately narrow approach* to the problem at hand. Once we set that inappropriately narrow approach aside and instead expand the scope of our both our factual and our theoretical inquiry – once we toss off the blinders – two things happen. First, facts begin to appear on our moral radar that we had previously ignored or considered irrelevant. And, second, principles that previously seemed to capture the person-based intuition in perfectly adequate ways come to seem inappropriately narrow.

Boonin's analysis, in effect, blinds us both to the morally relevant factual details of our own cases and to alternate formulations of the person-based intuition. A more expansive view is in order – if, that is, the person-based approach is to have any hope of success at all.

We may still find ourselves on Boonin's side in the case of Wilma and Pebbles. But we may at the same time be able to open the door to a finding of wrongdoing in those versions of the non-identity problem in which the

choice under scrutiny seems to us clearly to be wrong.¹⁵

2. BOONIN'S ARGUMENT

According to the person-based intuition, what is “bad” is “bad for” some existing or future person. Boonin unpacks the overall structure of the intuition in the form of premises (4) and (5) of the non-identity argument he constructs for the Wilma and Pebbles case: if an act harms no one, it wrongs no one; and unless an act wrongs someone or another, it's not wrong.¹⁶

I have no particular quibble with Boonin's work in that connection – that is, with Boonin's premises (4) and (5). For my purposes here, I will focus on how he unpacks not the overall structure of the intuition but just the *necessary condition* that the intuition sets forth. When, on Boonin's view, is that necessary condition satisfied? When does a choice *harm* a given person?

That question is answered by Boonin's premises (1) and (2). Here, I

¹⁵ I do agree with Boonin that Wilma's choice is permissible. But I want to underline that I do not accept Boonin's view that almost anything goes when it comes the burdensome, existence-inducing, choices that agents may make. Contrary to the position Boonin attributes to me, I have consistently rejected that position – including its application in cases involving commercial surrogacy and human cloning, with what I meant to be a vengeance. Boonin (2014): 14 n. 21; Boonin (2014): 189-290 n.1; cf. Roberts (2009), Roberts (2007), Roberts (1998), Roberts (1996).

It's true that I have defended a person-based form of maximizing consequentialism. Such an approach has seemed attractive to me in part in virtue of the support it provides to Narveson's famous dictum that “[w]e are in favour of making people happy” but “neutral about making happy people.” Narveson (1976): 73. (Narveson's position is helpful in answering the question whether, other things equal, we are obligated to bring ever more happy people into existence. Narveson thinks, and I think, we are not.) But my defense of the person-based approach largely hinges on a demonstration that, in the sorts of cases where it's very clear to us that what has been done is morally wrong, the reasoning behind the nonidentity problem itself indeed goes awry. On my view, Boonin's nonidentity argument, applied, e.g., to Parfit's Depletion and Risky Policy cases, Kavka's Slave Child and Pleasure Pill cases, cases involving historical injustices and, now, the Warren and Percy case, in fact fails.

Boonin is correct to include me in that very small class of philosophers who think that a person-based approach can be made to work. But I think it can be made to work because the approach can be articulated in a way that gives it the resources that it needs to avoid the result of permissibility in the nonidentity cases in which it seems very clear to us that what has been done is morally wrong. Thus Boonin is mistaken to include me in that still smaller class of philosophers who, when it comes to our burdensome, existence-inducing choices, and specifically those we can agree are clearly wrong, claim to accept the Implausible Conclusion.

¹⁶ Boonin (2014): 5.

generalize on the basis of the argument he makes in the Wilma and Pebbles case.

- (1) In the non-identity cases, making the choice under scrutiny *c* rather than an alternate choice *c'* doesn't make the person *p* – e.g., Pebbles – burdened by *c* worse off than *p* “would otherwise have been”; and
- (2) A choice *c* harms a person *p* only if *c* makes *p* worse off than *p* “would otherwise have been.”¹⁷

An assumption of the argument – an assumption that comes to us in the form of Boonin's premise (3) – insures that we need not worry about anyone other than the future person. We are thus to assume in the case of Wilma and Pebbles that neither Wilma nor any other existing or future person is made worse off, or harmed, by Wilma's choice to bring the blind Pebbles into existence rather than the sighted Rocks.

The combination of these five premises do indeed seem to generate the conclusion that Boonin suggests: that Wilma's choice is morally permissible.

As noted, my interest here is in how Boonin formulates the necessary condition the intuition is meant to establish.¹⁸ In part 4 below, I will argue that that formulation is inappropriately narrow.

The discussion there will focus on the conceptual component of the non-identity problem. But I'll start in part 3 below by taking a look at the factual component of the non-identity problem. Does Boonin leave some of the factual details of some of his own cases – details we can't assume in advance to have no moral relevance to our evaluation of the choices under scrutiny – unmentioned or unexplored?

¹⁷ Boonin (2014): 3-4.

¹⁸ I begin here – in part 3 – with the factual analysis. In part 4, I turn to the theoretical analysis. I will there argue that Boonin's understanding of how the person-based intuition's necessary condition operates – as reflected in his own formulation of that condition – is inappropriately narrow. It blinds itself to the morally critical factual details of the relevant cases. Perhaps it's that the failure to appreciate those details from the beginning leads Boonin to adopt a crude formulation of the condition. Or perhaps the explanation goes in the other direction: perhaps the problem is that the crude formulation never requires us to look closely at the morally relevant details of our own cases. No matter. If the person-based intuition is to be made plausible at all, then the necessary condition that it asserts must be formulated in a way that is sensitive to the morally relevant details of the various choices that it deems morally permissible.

3. THE NARROW VIEW OF THE FACTS OF THE NON-IDENTITY CASES

3.1 *Direct version.* Consider, again, the Warren and Percy case. In my original description of the case, I didn't explicitly note that, in addition to (1) the option of introducing the chemical into the dish and allowing fertilization to proceed and the new person, Percy, to come into existence and (2) the option of discarding the contents of the dish, Warren also had (3) the option of not introducing the chemical into the dish yet still allowing fertilization to proceed and Percy to come into existence. It's a stipulation of the case that, had Warren not chosen option (1), he *would* have chosen option (2). But it doesn't follow from that stipulation that option (3) is not a perfectly available alternate choice for Warren at the moment just prior to choice. Indeed, the most natural reading of the case is that that alternative existed for Warren but he simply declined to avail himself of it. Perhaps he declined because he then wouldn't have had the chance to challenge accepted wisdom as to just how dangerous the chemical really is; perhaps he wouldn't have gotten the \$100. But none of that even begins to suggest that that third alternative option did not exist for Warren. At least: we have no reason to think it doesn't. And it's a choice that would have made things better for Percy than things in fact are – as, that is, things are for Percy, given Warren's choice of (1).

This is not to suggest that Warren had the *obligation* to choose (3) rather than (1). Indeed, any plausible formulation of the person-based intuition will generate the result that Warren had no such obligation.¹⁹ It's just to say that the fact that, for all the reader is told, (3) is an alternative for Warren is a detail of the case. Moreover, it's a detail that we cannot assume in advance to be irrelevant to the evaluation of Warren's choice of (1).

Thus, there is a distinction between the Warren and Percy case and the Wilma and Pebbles case: in the Wilma and Pebbles case, as Boonin describes that case, it really doesn't seem that the agent has any third option, but in the Warren and Percy case, it seems the agent does have a third option. The one case is thus a two-option case, and the other a three-option case. We can't, in advance, take it for granted that that distinction itself does not mark a morally relevant distinction between the two cases.

3.2 *Indirect version.* As noted earlier, Boonin's presentation of the non-identity problem – in the form of the Wilma and Pebbles case – is clean and clear. The difficulty is that the many versions of the non-identity problem – those in which we want to say the choice is wrong but in which the most

19 See note 15 above.

plausible victim of the wrongdoing may seem not to have been made *worse off*, or *harmed*, at all – *aren't* clean and clear. Rather, they are complicated and a bit murky.

A good example of such a version of the non-identity problem is the Risky Policy case. Thus consider the “wealthy society” that is “running out of fossil fuels.”²⁰ That society faces the choice between two sources of energy, one that would allow existing people to retain their standard of living and have no effect on future people and one that would somewhat improve the standard of living for existing people but also generate a “significant amount of toxic waste,” an amount that would eventually “painlessly” kill “tens of thousands” of future people “once they reached the age of forty.”²¹ It’s stipulated that “over time, the effects of [the] subtle differences” in when and with whom people conceive their children “will be enough to generate two entirely distinct sets of people: the set of people who will exist five hundred years from now if the safe policy is selected, and the completely different [non-identical] set of people who will exist five hundred years from now if the risky policy is selected.”²² It’s thus a stipulation of the case that, if “the wealthy society had instead selected the safe policy, those innocent people would never have existed in the first place.”²³

On the basis of those facts, Boonin notes the following:

The No Worse Off Claim. The “innocent people who are killed as a result of the leaking toxic waste are not made worse off by the wealthy society’s choice of the risky policy” than they would have been had the wealthy society chosen the safe policy instead.²⁴

The No Worse Off claim is actually just premise (1) of the non-identity argument that Boonin wants to say is sound. But our interest in the No Worse Off Claim at this juncture is not to ask whether it’s true or not. Rather, what is of interest here is what the claim reveals regarding Boonin’s view of the facts of his own case – specifically, that Boonin’s view of the case is itself very narrow.

Consider any possible person – say, *Polly* – who is a member of the class of people who will come into existence and die painlessly at age forty if the risky policy is chosen and who would never have existed at all if the safe

20 Boonin (2014): 5.

21 Boonin (2014): 5-6.

22 Boonin (2014): 6.

23 Boonin (2014): 6.

24 Boonin (2014): 6. This assertion would fit the pattern Boonin provides for premise (1) of the non-identity argument. Boonin (2014): 3.

policy is chosen. We should agree with Boonin that it's better – or at least not worse – for Polly to come into existence and die painlessly at age forty than it is for Polly never to exist at all. And we should accept the counterfactual stipulation that, had the safe policy been selected instead of the risky policy, Polly *would* never have existed at all. And for the sake of clarity let's just suppose that the risky policy has in fact been selected and that Polly in fact *will* exist and will die painless at age forty.²⁵

The easiest way to see that Boonin's presentation of the facts of the case is at least arguably overly narrow is to reformulate what we have just said in the language of possible worlds. Thus: in the actual world – call it *w1* – in which the risky policy is in fact chosen, Polly eventually exists but is doomed to die painlessly at age forty, and in the closest possible world where the safe policy is chosen instead – call it *w2* – Polly never exists at all. Even given that it's better – or at least not worse – for Polly that Polly exist and die painlessly at age forty than it is for Polly never to have existed at all, all that we can infer from what we have said so far is that the choice of the risky policy *as that policy is put into effect at w1* doesn't make things worse for Polly than the choice of the safe policy *as that policy is put into effect at w2*. We can't, that is, on the basis of what we have said so far rule out that there exists a world – call it *w3* – in which the safe policy – or at least a *safer* policy – is itself implemented *in some alternate* way that both brings Polly into existence and makes Polly better off than Polly is under the risky policy as that policy is implemented in *w1*.

Boonin's own presentation of the case has us take a narrow view; it has us put our blinders on. We are to attend exclusively to the risky policy at *w1* and the safe policy at *w2*. We determine that the one is better – or at least as good – for Polly as the other.

But now let's remove the blinders. We then recognize that there's more than one way to implement the safe policy. The safe policy can also be implemented in the way it is implemented in *w3*.

This is not to add facts to Boonin's case. This is just how the world is. The "precariousness" of any particular person's coming into existence does not mean that there's *only one possible route* for that person ever to come into existence at all. In fact, the choice is made to adopt the risky policy – and as a result of that policy Polly eventually exists and dies painlessly at age forty. That is what happens at *w1* – at, that is, the *actual* world. But it *could* have happened – specifically, it's physically and

25 Boonin's language on this point goes back and forth between a simple material conditional and the corresponding counterfactual. For purposes of evaluation, this distinction doesn't seem material.

metaphysically possible – that the agents who made the one choice at w_1 have chosen the safe policy instead and then implemented that policy in a way brings Polly into existence and makes Polly better off than she is in the actual world. That possibility – that alternate outcome – is represented by w_3 .

And that's a factual detail that's left unmentioned but would seem to be implicit in the case. It's a *modal* detail, to be sure, involving what could have been. But it's still part of the case.

The upshot? While it's clearly true that the risky policy *at* w_1 makes things *better for Polly* (or at least not worse for her) than the safe policy *at* w_2 , it's also clearly true that the safe policy *at* w_3 makes things *still better* for Polly than the risky policy at w_1 .

But if that's so, then it would be a mistake – starting out – to assume in advance that that particular factual detail isn't relevant to our moral evaluation of the choice of the risky policy. It would be a mistake to assume in advance that we can complete the moral evaluation without taking that detail into account. Now, perhaps, consistent with all that we have said so far, that detail will turn out *not* to be critical to our evaluation of the case. My point here, however, is that it's a mistake to assume – starting out – that it's not relevant to our evaluation of the case.

Familiar “same people” cases help to underline that fact. Consider, for example, Medical Case I. Suppose that the doctor has the choice to treat you with medicine A, not to treat you at all or to treat you with medicine B. Suppose that the doctor chooses medicine A and that, while medicine A saves your life, the life you are left with is, though clearly worth living, substantially diminished. It's nice that the doctor has saved your life, given that (let's suppose) you'd have died had the doctor not treated you at all. But it's surely morally relevant to our evaluation of the doctor's choice to treat you with medicine A that medicine B would have *both* saved your life *and* restored you to full health.

This is not to contest the No Worse Off Claim as Boonin has articulated that claim in connection with Risky Policy. The claim is perfectly true.²⁶ The risky policy at w_1 is better for Polly – or at least not worse for Polly – than the safe policy at w_2 . But the claim that the safe policy at w_3 is better for Polly than the risky policy at w_1 is also perfectly true. Any assumption that the first detail regarding the case but not the second will be morally

26 The language of possible worlds certainly helps to avoid any quibbles in this context. In contrast, any implication that as a matter of physical necessity p can't exist under the choice of the safe policy – that in all physically accessible possible worlds the choice of the safe policy, or at least a safer policy, results in p 's never existing at all – would be highly problematic.

relevant to our evaluation of the choice of the risky policy would seem to be – and clearly is, on my view – illegitimate. Of course, a critic may argue that I have made things too complicated and that a stipulation or two will clean things up in a way that shows that we should have no objection to Boonin’s – or anyone else’s – narrow view regarding of the facts of Risky Policy. Indeed a number of superficially appealing strategies for justifying the narrow view are available. I will consider three such strategies here.

(i) *Counterfactual stipulation*. It might be argued that my insistence on a more detailed description of Risky Policy ignores a main stipulation of the case: that had the risky policy not been chosen, then the safe policy would have been selected in its place and Polly would never have existed at all. We can articulate that stipulation in the language of possible worlds. The risky policy is implemented in the actual world w_1 and Polly eventually exists there. Moreover, at the closest possible world w_2 at which the risky policy isn’t implemented and the safe – or a safer – policy is implemented instead, that policy is implemented in such a way as to leave Polly out of existence altogether.

The counterfactual stipulation may seem to provide a rationale for Boonin’s narrow view of the case. It may seem to ground the position that the fact that the safe policy as that policy is implemented at w_3 is irrelevant to the moral analysis. In fact, however, counterfactuals can seem morally telling when they aren’t. In Medical Case I, the doctor who treated you with medicine A may offer the defense that he has saved your life. And indeed he did. He may even convince us that the relevant counterfactual holds: that is, that he would have declined to treat you altogether (perhaps exiting the hospital in a moment of pique) and you would have died had he not treated you with medicine A. We might (momentarily) be duped into thinking that the doctor is, after all, both morally and legally off the hook.

But he’s not. What he’s said in his defense is perfectly consistent with the further point that treating you with medicine B would have restored you to full health and thus left you better off than you in fact are. That that option exists is a function of what the *world* – metaphysically, physically – is like; it’s not a function of what the *doctor counterfactually* would or would not choose to do. It’s stipulated that treating you with medicine B isn’t an option the doctor *would* have chosen. But it doesn’t follow from that that it’s not an option available to the doctor. The world, not the doctor, determines what the doctor’s options are.

We clearly think that the factual detail involving the option of treating with medicine B has moral salience in Medical Case I. There’s no basis for concluding that the analogous factual detail in Risky Policy – the detail involving the many different ways in which the safe policy might itself be

implemented; the detail that makes it so that w3 is an alternate possible way in which the future might have unfolded – doesn't have moral salience in that case.

(ii) *Stipulation of physical necessity.* Alternatively, Boonin might have meant his reader to understand that it's a stipulation of the case that, as a matter of the physical law, the safe policy can be implemented in only one way, a way that inevitably and unavoidably leads not to w3 but rather to w2, where Polly never exists at all.

It's true we could devise a case in which the stipulation of physical necessity holds.²⁷ For example, we could add to the case that the toxic waste produced under the risky policy in w1 – exactly that degree of risk and exactly that degree of toxicity – is physically *necessary* to Polly's ever coming into existence at all given her genetic makeup.

Now, Boonin doesn't seem to have that particular story line in mind – nor did Parfit, in his presentation of his own version of Risky Policy. Unlike the counterfactual stipulation, the stipulation of physical necessity seems completely foreign to the case that Boonin presents. It's not one the reader can be expected to bring to the table. If it's to be included in the case at all, it should be included *explicitly*.

But if we drop the stipulation, we shall need to accept that the choice of the safe policy can, after all, be implemented in more than one way – that is, that as of that moment just prior to choice the future isn't carved out between two and only two paths forward. That's just the way the actual world happens to be, given the physical laws that in fact govern that world. The safe policy doesn't *as a matter of physical necessity* mean that Polly *can't* exist. We dwell in more possibility than *that*. Dropping the stipulation in effect returns us to the earlier account of the case sketched above.

Of course, nothing forces us to drop the stipulation. That is: we can always construct another version of the case, a version that – explicitly – includes the stipulation and thereby strips w3 out of the picture altogether.

It seems clear, however, that that revised case is very different from the case Boonin starts us off with. In Boonin's case, situated (sans stipulation) in the world that is in important respects very much like the actual world,

27 It's possible Boonin himself intends Risky Policy to include the stipulation of physical necessity. Moreover, as a referee for this paper points out, "there are no facts to the case" beyond the facts the author of the hypothetical describes; it's Boonin's case to detail as he pleases. At the same time, Boonin doesn't explicitly include the stipulation in his description of the case. And the case where the stipulation isn't included and the case where the stipulation is included are two distinct cases. In this part 3.3.2, I am considering how the case plays out both ways – first, in this subpart (ii), as it plays out with the stipulation in place and, second, in subpart (iii) below, as it plays out without the stipulation in place.

we think it's clearly wrong for agents to choose the risky policy. But in the revised case it's not at all clear the choice is wrong. In the revised case, the risky policy is more like a fertility treatment, a treatment that produces a certain adverse side effect for Polly but is physically necessary to Polly's ever coming into existence at all. (Perhaps the genetic materials that eventually develop into Polly are such that cell division in the early days after conception *cannot* happen in the absence of the toxic waste; or perhaps the physical laws that govern the actual world simply don't apply.) In that case, Polly *can't* exist and *not* die painlessly at forty, just as I can't exist and not die, I hope painlessly, within the next few decades.

It's just not clear that on those facts – all known to the agents and all matters in respect of which agents have certainty – the choice of the risky policy is wrong. After all, we accept that the choice of a fertility treatment, even one that has certain adverse side effects, may be morally permissible. We may also accept that the choice of the risky policy – made in some distant possible world in which human reproduction physically *requires* exposure to toxic waste in a way that in the actual world human reproduction physically *requires* exposure to certain hormones – is permissible as well.

(iii) *Relevance of probabilities.* A third strategy for justifying Boonin's narrow view regarding the facts of Risky Policy concedes that, technically, w_3 exists as an available option but at the same time insists that the *probability* that w_3 will unfold – that is, that Polly will exist – given the choice of the safe policy is *extraordinarily* low and thus can, and should, be disregarded altogether for purposes of evaluation.²⁸

That probabilities may be relevant to how we evaluate choice is demonstrated in Medical Case II. In that case, the doctor is deciding between medicines A and B. The doctor is fully informed, let's suppose, as to the following facts. First, the probability that medicine A will restore the patient to one degree of good health or another within a relatively narrow range – as opposed, let's say, to killing the patient – *is truly significant*. Whatever it is, it's a probability that the doctor should take into account in making his choice. Second, the probability that medicine B will restore the patient to some *greater* degree of good health but still within that same narrow range is also truly significant. It, too, is a probability the doctor should take into account in making his choice. And third, the probability that medicine C will restore the patient to an even greater degree of good health but still within that same narrow range is *very low, practically nil*. In

28 As one referee for this paper points out, the idea that Boonin's evaluation of the choice in Risky Policy turns on the relevant probabilities "comes rather out of the blue." That's so. The effort here, then, should be understood as an effort to understand how Boonin's argument might be reconstructed.

those circumstances, certainly the doctor needs to compare the different outcomes made possible by medicine A on the one hand and medicine B on the other, along with the probabilities that accompany each, in order to make his decision. But does he really need also to compare A against C and B against C, given the probabilities? No. Rather, it seems perfectly appropriate in that case for the doctor to decide what to do on the basis of the comparison of A against B and on that basis alone.

Now, for purposes of completing that comparison, it will be important to have detailed information regarding the relevant possible outcomes and the relevant probabilities. If the outcome under medicine A would be *better* for the patient than the outcome under medicine B, and if the (significant) probability of success under medicine A is also *greater* than the probability of success under medicine B, then surely medicine A is the correct choice.

We should, in other words, clearly recognize that probabilities may well be relevant to moral evaluation. However, what is often forgotten in discussions of the non-identity problem is that it's notoriously easy to make mistakes in calculating the relevant probabilities.

Setting C aside in Medical Case II may make sense, given that the probability of success for the patient under C is – by hypothesis – very low and practically nil. It is much lower than the probabilities of success for the patient under A and B which are – by hypothesis – significant. *Probability differentiation* is, in other words, in effect a stipulation of the case.

But Risky Policy is very different. We don't see – at least, we are making things up if we think we *do* see – any analogous probability differentiation in the facts of Risky Policy.

Why is that? Isn't it just *part* of Risky Policy that the probability that Polly will exist under the choice of the risky policy is significant while the probability that Polly will exist under the choice of the safe policy is very low, practically nil?

The latter issue – just what the subject's chances of existence are under the safe policy – is occasionally explicitly resolved in various descriptions of various non-identity cases.²⁹ But the former issue – just what the subject's chances of existence are under the risky policy – almost never is. Rather, what is typically included in those descriptions is just that, given the choice under scrutiny, a particular person (e.g. Polly) *in fact* comes into existence and *in fact* suffers in a certain way. The question of the *probability* that Polly will exist under the choice of the risky policy is never even raised,

29 See Kavka (1982): 100 n. 15.

much less answered.

So let's raise it now. Consider Polly, who, by hypothesis, exists and suffers as a result of the choice of the risky policy. What were Polly's chances of coming into existence, calculated at the moment just prior to that choice, given the choice of the risky policy? To think that Polly's chances of coming into existence, as of that moment just prior to choice, were anything other than very low, practical nil – to think that they were *significant* – is to fail to appreciate all the many ways that agents could have, as of that moment just prior to choice, gone about implementing the safe policy. Just as the fact that there existed many, *many* ways of implementing the safe policy just prior to choice meant that Polly's chances of coming into existence under that choice are very low, so does the fact that there existed many, *many* ways of implementing the risky policy just prior to choice mean that Polly's chances of coming into existence under that choice were very low.

To think otherwise is to think of the risky policy as a sort of fertility treatment that Polly's forebears might have unwittingly undergone – to think of the risky policy as having a sort of existence-inducing *power* in respect of the particular person Polly's coming into existence.

Can't we just stipulate that that's so? Of course. But then what we are calling "Risky Policy" is no longer Risky Policy. "Risky Policy" is now a fertility treatment case. And we do accept some adverse side effect, in connection with fertility treatments, when the treatment itself makes it more likely that a couple will conceive a child whom they badly want. And we evaluate fertility treatment cases in light of that fact. But Risky Policy in its original form is *not* a fertility treatment case.³⁰

I noted earlier that philosophers' descriptions of the various non-identity cases typically include the fact that, the choice under scrutiny having been made, a particular person, e.g. Polly, in fact exists and suffers. In that connection, it's important to underline that it's a mistake to judge how *probable* an outcome is given a particular choice on the basis of how the future *in fact* unfolds. Rather, the calculation of probability, for purposes of the moral evaluation, should be based on information available to the agents just prior to choice, the whole point of that calculation being to provide the agents with a means, prior to choice, for determining

30 When it comes, e.g., to cases involving historical injustice, it may be that the choice we all agree is wrong – e.g., the choice to order or to participate in the implementation or continuation of the Holocaust – in fact makes it *less* likely that particular offspring of the contemporaneous victims of that wrong choice will ever come into existence at all. The same point applies to at least some other versions of the non-identity problem, including Depletion and Risky Policy.

permissibility.³¹

My argument, then, is that there is no basis for thinking that the choice of the risky policy, as of that moment just prior to choice, generates any significant probability of Polly's coming into existence at all – and certainly no greater a probability than the choice of the safe policy would have generated for Polly. For both choices, there are just too many ways of implementing that choice – the risky policy on the one hand, the safe policy on the other – to give Polly *any significant chance of existing at all*.³²

31 An assessment of the probabilities typically is brought to bear to explain why the choice that looks to be the morally correct choice just prior to choice can't subsequently be dismissed as morally wrong when the future unfolds in some way that no one could have predicted. (Suppose by some quirk the seatbelt I've asked you to put on ends up strangling you when we are hit by a Mack truck speeding down the highway in the wrong direction; we don't think on those facts my asking you to put your seatbelt on is morally wrong.) Or why agents are not thought to have done anything wrong in the case where a future that would have been better for a particular person existed as a physical possibility for the agents but the agents has no clue how to bring that physical possibility about. (I *can* save your life by performing open heart surgery on the subway, but I have no idea how to go about that; we think on those facts my instead calling 911 and performing CPR is a perfectly acceptable alternative.) An assessment of probabilities, in other words, typically arises once it's determined that an expected value calculation is in order. That means, in turn, that the probabilities we should be focusing on are those we calculate (1) just prior to choice – at no later or earlier point – and (2) just on the basis of information available to the agents at that moment.

32 A referee for this paper suggests that my line of reasoning in this connection may be challenged by an alternate way of coming at the calculation of the probabilities at issue. According to that alternate approach, we first recognize that some person or another will exist (and suffer) under the choice of the risky policy. We then simply ask how likely it is that that very same person, "whoever it is," would have existed under the choice of the safe policy. Of course, the answer to that question is that the chances of that person's existing under the safe policy are very low. The suggestion, as I understand the referee's comment, is that we are then to compare the fact of that person's (burdened) existence under the choice of the risky policy against the very low probability that that same person would exist (and have a relatively unburdened existence) under the choice of the safe policy. It may very well be that many philosophers have – assuming they have contemplated the various probabilities at all – understood the analysis to proceed in just that way. Stepping back, however, I think we can see that on this approach what we are asked to compare is the value the future that *in fact* unfolds for that person – arbitrarily, *p* – given the choice of the risky policy against how the future can be *expected* to unfold for that same person *p* given the choice of the safe policy. We are asked to compare, that is, an *actual* wellbeing level against an *expected* wellbeing level – a *ex post fact* against an *ex ante* projection. But the principle that would deem such comparisons relevant for determining the choice that is better for the subject will lead to inconsistency, potentially directing in, e.g., a medical case in which the patient's chances of survival are low whatever choice is made, that medicine A is better for the patient than medicine B (A producing more actual value for the patient than B produces expected value) and that medicine B is better for the patient than medicine A (B producing more actual value for the patient than A produces expected value). I take no position here on whether an actual value calculation of wellbeing or an expected value calculation is correct for purposes of moral evaluation. But it does seem clear that we can't mix the two. For further discussion, see Roberts (2009); Roberts (2007).

But that means that the probabilities don't, after all, provide a basis for Boonin's narrow view. We can't disregard the third option on the ground that it's highly improbable that Polly will ever exist at all under the choice of the safe policy. For it's just as highly improbable – as of that moment just prior to choice – that Polly will ever exist under the choice of the risky policy. We thus can't assume, on the basis of the probabilities, that the third option won't have relevance for purposes of our evaluation of the agents' choice of the risky policy.³³

4. THE NARROW VIEW OF THE CONCEPTUAL COMPONENT OF THE NON-IDENTITY PROBLEM

We now turn to the conceptual component of the non-identity problem – in particular, how we should formulate the person-based intuition's necessary condition on wrongdoing. My argument here will be that Boonin's narrow view regarding how that condition is to be formulated is a mistake.

4.1 Concept of harm. Let's begin with a close look at Boonin's premise (2). Generalized:

(2) A choice *c* harms a person *p* only if *c* makes *p* worse off than *p* "would otherwise have been" had *c* not been chosen.³⁴

Boonin's reference to the concept of harm in his formulation of the person-based intuition seems to me to be unproblematic³⁵ – as long as we get the concept of *harm* right. Thus: the intuition itself would be that a choice is wrong only if it *harms* a person who does or will exist given that choice.³⁶

Moreover, the comparative account of harm that Boonin proposes also seems to be on the right track. Pain itself, e.g., isn't enough to show harm (in, as Parfit put it, a morally relevant sense). A vaccination may inflict *pain* but shouldn't *harm*. Harm requires, rather, that the person harmed has been made *worse off*.

33 I have made a similar argument elsewhere. See Roberts (2009); Roberts (2007).

34 Boonin (2014), pp. 3-4. Boonin, referring to a "commonsense understanding" of harm, words the point this way: "If your act harms someone, then it makes that person worse off than they would have been had you not done the act." Boonin (2014): 3.

35 This we find in his premise (4).

36 Here, I follow Boonin in sharply distinguishing harming from wrongdoing. Harming is a matter of (comparatively) reducing a person's wellbeing – of a person's having less wellbeing in one possible future, or outcome, than in another. On that view, a person may be harmed without having been wronged. Philosophers who want to understand harm itself to come with clear normative implications may want to understand my "harm," and Boonin's, to constitute something like a comparative loss of wellbeing rather than a bona fide harm.

The difficulty for Boonin is that the particular comparative account of harm that he offers seems overly narrow. It compares how things have unfolded for a particular person under a particular choice against how things would have unfolded had that choice not been made. It elevates, in other words, the counterfactual stipulation – in, e.g., the Warren and Percy case and the Risky Policy case – to the status of primary conceptual significance. And that’s surely a mistake.

Thus Boonin writes: “What seems to be a commonsense understanding of what is required in order for an act to harm someone” is the following:

if your act harms someone, then it makes that person worse off than they would have been *had you not done the act*. That this is a widely accepted view seems to be confirmed by the fact that if a person is accused of harming someone, it is standardly taken as a sufficient rebuttal to the claim if they can establish that they have not made their alleged victim worse off than they would otherwise have been [“had you not done the act”].³⁷

To see that this account of harm won’t do, let’s go back to Medical Case I.³⁸ It’s part of that case that, had the doctor not treated you with medicine A, he’d not have treated you at all (perhaps exiting the hospital in a moment of pique). His defense is that he has not made you, the alleged victim, worse off than you “otherwise would have been” – indeed, that he’s saved you

37 Boonin (2014): 3. Emphasis, and square-bracketed material, added.

38 Boonin’s account of harm is thus very different from my own. Thus Boonin’s account is unable to identify harm in either Medical Case I or the Warren and Percy case. In both those cases (as a referee for this paper has pointed out) the agent in a sense makes it the case that the impaired child must suffer the impairment or never exist at all. On Boonin’s account, that is enough to show that the child isn’t harmed. On my account, that the agent would decline to make a third choice that would have made things still better for the child is irrelevant to the analysis – a red herring. It’s what the agent could have done – physically, metaphysically – that is salient to the case. The world, not the agent, determines what the agent’s options are. In both these cases, that that third choice existed as an alternative for the agent shows that the one choice harms the child.

Of course, it’s also part of the person-based approach that it’s permissible never to bring the child into existence to begin with. Perhaps that choice itself can appropriately be said to harm the child or at least to cause the child to sustain a loss of wellbeing. I would agree – but also say that, under a person-based approach, the sorts of harms or losses we sustain by virtue of our never having existed are morally without any significance at all. The principle that losses should be distinguished in that way I call the Loss Distinction Thesis (or, in earlier work, Variabilism). Roberts (2011a), (2011b).

It should be noted that cases involving causal over-determination and preemption challenge Boonin’s account of harm as well. While I don’t survey such cases here, any plausible account should be able to identify harm in those cases. Thus on my view the bare fact that a group of agents had the option of creating more wellbeing for a given person *p* may be enough to show, depending on the details of the case, that the choice of any one member of that group by virtue of his or her “participation” in the group harms *p*. Roberts (2006).

life. And it's perfectly true that had he not "done the act" – had he not treated you with medicine A – you'd have died. Thus, he has not *harmed you*.

But how plausible is the doctor's defense? *Moral law aside, human law* would certainly scoff at such a defense and insist rather that a mistake has been made in identifying the appropriate *baseline* (a concept Boonin, and before him Feinberg, freely relies on). For purposes of the legal analysis, the baseline isn't what the doctor *would* have done to you – that is, still worse – had he not done what he did. Rather, it's what the doctor *could* have done for you given the options available to him at the critical time. Thus, the difference between your diminished life under medicine A, and the life you would have had under medicine B, a treatment that would have restored you to full health, provides the appropriate measure of harm in Medical Case I.

It's true that the choice under which you have the diminished life is better for you than the choice in which you simply die. But that truth is just a red herring: it tells us nothing about the harm that is done to you by the choice to treat with medicine A at all.

The case of Warren and Percy makes the same point. It's a stipulation of the case that had Warren not introduced the chemical into the dish, thereby causing Percy to be born with an impairment, Warren would have simply tossed the contents of the dish. What Warren in fact does nonetheless *harms* Percy. And what tells us that that's so is that Warren has a third option: he could have refrained from introducing the chemical into the dish and nonetheless allowed fertilization to proceed and Percy to be born. If we fail to attend to that factual detail of the case – or if our formulation of the relevant person-based necessary condition on wrongdoing sets that particular factual detail of the case aside as irrelevant to the moral evaluation – then that evaluation itself will be unreliable.

There's no reason to think that that more expansive concept of harm isn't the correct concept for purposes of formulating our person-based principles as well. Yes, the principles will be a (tiny) bit more intricate. But that's a small price to pay to avoid a principle that generates clearly false results across a wide range of cases.

The proposal here, then, is that Boonin's overly narrow analysis of harm – what we see in his premise (2) – be rejected, in favor of a more expansive understanding of harm. Thus:

Expansive premise (2): A choice *c* at a world *w* harms a person *p* only if there is some choice *c'* that exists as an option for the agent (or agents) and some alternate accessible world *w'* such that *c* at *w* is worse for *p*

than c' at w' is.

That is: a choice harms a person p only if that choice makes things worse for p than at least some available alternate choice makes things for p – only if, that is, agents could have done more for p than they in fact have. If c is maximizing for p , in other words, it doesn't harm p .

Now, moving from harming p to wronging p – or to wrongdoing – isn't an easy task; all kinds of tradeoffs need to be taken into account. But at least on this conception of harm, the door is wide open to a finding of harm – which means the person-based intuition is wide open to a finding of wrongdoing – in all the right cases: in, e.g., the Warren and Pebbles case and the Risky Policy case. We still find ourselves in the controversial position – with Boonin – of declaring Wilma's choice perfectly permissible. But perhaps that's not such a bad place to be. (How could we live without Pebbles?)

4.2 *The red herring*. Let's now turn to Boonin's premise (1):

- (1) In the non-identity cases, making the choice under scrutiny c rather than an alternate choice c' doesn't make the person p – e.g., Pebbles – burdened by c *worse off* than p “would otherwise have been” under c' .

And let's just note that, once we restore the counterfactual stipulation itself to its rightful place – once we focus, not on what the agents would have done, had they not done what they in fact did, but rather on what agents could have done instead – we can see that the comparison that premise (1) asks us to make – the comparison that launches Boonin's own version of the non-identity problem – may be (depending on the factual details of the particular case) just a red herring.

Let's step back. In the Wilma and Pebbles case, the comparison of how things have turned out for Pebbles and how they would have turned out had Wilma not chosen as she did is – if misleading – relevant. It's relevant, in virtue of the fact that in that case Wilma just doesn't have the option of making things still better for Pebbles than they are under Wilma's choice not to delay the pregnancy. As argued in part 2 above, we need to *say* that: we can't *assume* that no further alternatives exist, or that if they do exist they aren't relevant to the moral evaluation. But in this particular case that happens to be so.

But in cases where still better alternatives do lurk – such as in the Warren and Percy case and in the Risky Policy case – (1) itself is just a red herring. It's always nice to hear that one option is better for a person than another, but until we know whether there's a better option still we are in no position to evaluate the one option.

5. CONCLUSION

I have argued that Boonin has been mistaken to take a narrow view of the cases that give rise to the non-identity problem. That narrow view causes him to overlook certain factual details – including critically important *modal* details – that are part of the cases that he himself develops. It may have led him as well to think that how we formulate the person-based necessary condition on wrongdoing is a simpler matter than it in fact is.

A way of summing up my argument is to say that Boonin's treatment of the most challenging non-identity cases – those (unlike the Wilma and Pebbles case) in which the choice under scrutiny is clearly wrong – mistakenly describes those *three*-option cases as *two*-option cases. He then – mistakenly again, in my view – tailors his principles accordingly.

While I have distanced myself from the conclusions Boonin reaches in many of those cases – and in particular from the Implausible Conclusion – I should underline again that I have a great deal more sympathy for Boonin's position as it applies in the bona fide two-option case, including, e.g., the Wilma and Pebbles case. Such cases aren't really so rare. Imagine a couple who decides to proceed with a pregnancy even though genetic testing has revealed that the resulting child will be impaired in some way; and imagine, as well, that the case is one in which no one other than the very child whose existence is at stake will be affected by the couple's choice. When such a couple decides to proceed, I am – like Boonin – committed to the position that the choice is morally permissible.

BIBLIOGRAPHY

- Boonin, D., 2014: *The Non-Identity Problem and the Ethics of Future People*, Oxford: Oxford University Press.
- Kavka, G., 1981: "The Paradox of Future Individuals," *Philosophy & Public Affairs* 11: 93-112
- Narveson, J., 1976: "Moral Problems of Population," in *Ethics and Population*, ed. Michael D. Bayles, 59-80, Cambridge, Mass.: Schenkman Publishing Company.
- Parfit, D., 1987: *Reasons and Persons*, Oxford: Oxford University Press

- Roberts, M., 1998: *Child Versus Childmaker: Future Persons and Present Duties in Ethics and the Law*, Lanham, Maryland: Rowman and Littlefield.
- 1996: “Human Cloning: A Case of No Harm Done?” *Jo. of Medicine and Philosophy* 21.5: 537-54.
- 2006: “Supernumerary Pregnancy, the Problem of Collective Harm and the Nonidentity Problem,” *Journal of Law, Medicine & Ethics* 34.4: 776-792.
- 2007: “The Nonidentity Fallacy: Harm, Probability and Another Look at Parfit’s Depletion Example,” *Utilitas* 19: 267-311.
- 2009: “The Nonidentity Problem and the Two Envelope Problem,” in *Harming Future Persons*, eds. M. Roberts and D. Wasserman, 201-228, Dordrecht: Springer..
- Roberts, M., 2011a: “The Asymmetry: A Solution,” *Theoria* 77: 333-367.
- Roberts, M., 2011b: “An Asymmetry in the Ethics of Procreation,” *Philosophy Compass* 6/11: 765-776.