



# The Subject of Harm in Non-Identity Cases

Jens Johansson<sup>1</sup>

Accepted: 26 September 2018 / Published online: 2 October 2018  
© The Author(s) 2018

## Abstract

In a typical non-identity case, the agent performs an action that causes someone to exist at a low but positive level of well-being, although an alternative was to create another, much happier person instead. There seem to be strong moral reasons against what the agent does, but it is difficult to explain how this can be so. In particular, it seems that on a simple counterfactual account of harm, the action does not harm anyone, as it does not make anyone worse off than he or she would have been had the action not been performed. In this paper, I criticize the response that the action nevertheless harms the actual child, as well as the response that it makes the counterfactual child worse off. I argue that the following alternative view deserves attention: a *compound* of the actual and the counterfactual child is made worse off by the action.

**Keywords** Non-identity problem · Harm · Causal account of harm · Counterfactual account of harm · Well-being · Merely possible individuals

## 1 Introduction

The non-identity problem arises in cases in which numerically different people will exist depending on which action is performed. Here is a typical case, borrowing most details from the main case in David Boonin's recent, book-length treatment of the non-identity problem (Boonin 2014; see also Parfit 1984: 358):

Wilma learns that if she conceives right away, her child is going to be born with an incurable disease. The disease would render his life just barely worth living. Wilma also learns that if she takes a pill each day for two months before conceiving, then her child will be perfectly healthy and have a much happier life. Wilma considers that option a bit too inconvenient, and conceives right away.

Intuitively, Wilma's action is morally wrong—at least assuming, as we will do, that it does not have any significant effects on anyone who is distinct from her child (such as Wilma herself, or

---

✉ Jens Johansson  
jens.johansson@filosofi.uu.se

<sup>1</sup> Department of Philosophy, Uppsala University, Box 627, 75126 Uppsala, Sweden

her child's school teachers or grandparents). It is no easy task, however, to explain *why* Wilma's action is wrong, or even to defend the claim that it *is* wrong. This would be fairly easy if Wilma's actual child—call him “Will”—were the same individual as the happier child—“Liam,” say—that she would have conceived had she waited for two months. Then we could say that Wilma's action is wrong at least partially in virtue of being harmful for this child, by depriving him of a life that would have been much better for him. The problem, though, is that her actual child and her counterfactual child are not plausibly the same individual. (While it might be metaphysically possible that Wilma conceives Will although she waits for two months—maybe because there is a possible world in which the very same sperm and egg that actually gave rise to Will unite two months later—it is unproblematic to suppose that this is not what would in fact have happened.) Since Will owes his (again, overall positive) existence to Wilma's action, it appears that he has nothing to complain about; nor, of course, does nonexistent Liam.

We can clarify and sharpen the problem by considering an argument—which, too, is primarily taken from Boonin (although similar lines of thought can also be found elsewhere)—for the claim that Wilma's action is morally permissible, an argument that appeals crucially to the apparent *harmlessness* of the action.

Premise 1: Wilma's action does not make anyone worse off than he or she would have otherwise been (that is, than he or she would have been had the action not been performed).

Premise 2: If Wilma's action does not make anyone worse off than he or she would have otherwise been, then it does not harm anyone.

Premise 3: If Wilma's action does not harm anyone, then it does not wrong anyone.

Premise 4: If Wilma's action does not wrong anyone, then it is morally permissible.

Conclusion: Wilma's action is morally permissible.

I have formulated all premises as claims specifically about the case at hand, instead of following Boonin in formulating the last three premises as general principles. As Molly Gardner has pointed out, stating the argument in the latter way would raise some irrelevant issues (Gardner 2015b).<sup>1</sup> For example, it would be unwise to formulate the third premise as the general claim that no action can wrong a person without harming anyone. Arguably, an action that does not harm anyone, and indeed makes no difference in any way to anyone's well-being, can still wrong a person—say, by violating the person's rights. Fortunately for Premise 3 as stated here, however, even if some factors independent of harming and well-being are sufficient for wronging, those factors do not seem to be present in the particular case of Wilma; for instance, it has not proven a successful strategy to say that her action violates Will's rights.

Indeed, even disregarding factors not related to well-being, if the second and third premises were put in completely general terms, their combination would be indefensible. This is worth noting, as it will lead us to an important stipulation about the Wilma case. Even ignoring non-well-being-related factors, it does not hold quite generally that if an action does not make anyone worse off than they would have otherwise been, then it does not wrong anyone. Suppose, for instance, that my action makes someone very badly off; and that if I had not performed this action, then I would have made him even worse off; and that yet another of my alternatives, which would have been easy for me to perform, was to simply leave him alone.

<sup>1</sup> As Gardner says, Boonin's way of stating the argument forces him to say, somewhat awkwardly, that even given certain sorts of counterexamples to a premise, the premise remains “true enough.”

Surely my action can wrong the person in such a case. It is debatable whether this sort of example involves harmful wrongdoing—the most reasonable position, in my view (Johansson and Risberg 2017)—or harmless wrongdoing: either way, the conjunction of the generalized versions of Premises 2 and 3 is refuted. By contrast, our specifically Wilma-related version of the argument can be immunized from this kind of criticism by stipulating—and I suggest we hereby do stipulate—that Wilma had only two alternatives: to create Will now, or Liam later (and that if she had not acted as she did, then she would have performed the alternative action).

The non-identity problem, applied to Wilma's case, is the challenge of either explaining why her action is morally wrong, in spite of the Booninian argument, or showing that the argument's conclusion, despite initial appearances, is acceptable after all. Boonin and a few others adopt the latter strategy (e.g., Boonin 2014; Schwartz 1978).<sup>2</sup> The emphasis is on "few," however, probably in part because the non-identity problem also arises on a larger scale, such as in issues of environmental policy, where the corresponding conclusion appears even more unappealing (Parfit 1984: 361–64). A more popular approach is to argue that Wilma's action is wrong not because of what it does to any particular individual or individuals, but, for instance, because it brings about less overall well-being than creating Liam would do. This strategy, which recommends denial of Premise 4, is often called—slightly misleadingly—the "impersonal" approach. While allowing us to resist the conclusion of the Booninian argument, the impersonal approach threatens to commit us to the dreaded claim that we have a strong moral reason, or even a moral obligation, to create happy people even when (as in the typical case) we can instead refrain from creating anyone at all, and to other implications that are often regarded as even more implausible than the conclusion of the Booninian argument (Boonin 2014: ch. 6).

I suspect that, despite their problems, one of these two approaches—Boonin's or the impersonal approach—is correct. There are, however, also various other strategies in the literature. Perhaps the most popular of these is the *harm-based* approach (e.g., Algander 2013; Bontly 2016; Gardner 2015a; Harman 2004, 2009; Shiffrin 1999). On this view, Wilma's action *is* harmful after all, and it is largely for this reason that her action is morally objectionable. This is the view that I aim to investigate in this article. My aim is significantly restricted, however, as I believe many versions of the view have already been adequately dealt with elsewhere (e.g., in Boonin 2014; Gardner 2015a, c). What I wish to do is, first, to argue that Molly Gardner's novel version, which despite its advantages over previous versions has *not* been dealt with already, is unsatisfactory (Section 2); and second, to consider, after some preparation (Section 3), the prospects of a completely different, and hitherto seriously underdeveloped version of the harm-based approach (Sections 4 and 5).

Note that merely arguing that Wilma's action harms someone would not amount to a harm-based *approach* to the non-identity *problem* (although it would, if successful, refute the Booninian argument as stated). Something needs to be added about the moral significance of this supposed harm. After all, if the harm did not make any moral difference, we could simply replace Premises 2 and 3 with the claim that if Wilma's action makes no one worse off than he or she would have otherwise been, then it does not wrong anyone. The resulting argument would not appeal to considerations of harm at all, but would still need to be addressed. A suitably ambitious version of the harm-based approach should say, then, that Wilma's action's being harmful does at least a substantial part of the job in explaining its moral wrongness.

<sup>2</sup> Heyd (1992) argues that identity-affecting cases like Wilma's are outside of the realm of morality, and that her act is thus not morally wrong (or, of course, morally permissible).

Some writers indicate that the non-identity problem should, instead, be taken to be the challenge of explaining why there is, or demonstrating the reasonableness of denying that there is, a moral *reason*—possibly outweighed by other considerations—against an action like Wilma’s (so that the conclusion of the Booninian argument is too weak to reflect the main problem). I am inclined to think that this is due to a failure to make the details of the case sufficiently clear. Admittedly, if Wilma’s action had meant, for instance, that some third party was spared a serious rights violation that she would have somehow committed by instead creating Liam, then it would be doubtful whether her action is morally wrong. However, given our stipulation at the outset that the case involves no such factors, I believe that most of us will have not only the intuition that there is a moral reason against her action, but also the intuition that her action is morally wrong. So I will primarily focus on the wrongness claim. However, it will be easy to see (although I shall not always make it explicit) that my main points also apply to the weaker claim that the harmfulness of Wilma’s action helps to establish a moral reason against it.

## 2 Harming the Actual Child: The Will-Focused View

On the standard version of the harm-based approach, Wilma’s action is morally wrong largely because it harms her actual child, Will. Call this the “Will-focused view.”

Since Wilma’s action does not make Will worse off than he would have otherwise been, the Will-focused view involves rejection of Premise 2—the claim that if Wilma’s action does not make anyone worse off than he or she would have otherwise been, then it does not harm anyone. As already indicated, in order to refute this premise it is not sufficient to refute the general counterfactual principle that an action harms someone only if it makes the person worse off than he or she would have otherwise been. Instead, proponents of the Will-focused view need to provide some sufficient condition on an action’s harming someone, which is satisfied by Wilma’s action and Will, respectively. Again, the focus here will be on Molly Gardner’s recent proposal (Gardner 2015a; see also 2015c, 2016, 2017a, b), but our discussion of it will eventually reveal a much more general problem for the Will-focused view.

Here is Gardner’s full account, which she calls the “Existence Account of Harming”:

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

*Harm*: 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. (Gardner 2015a: 434)

According to Gardner, the Existence Account deals adequately with the non-identity problem. Applied to the Wilma case, the thought is that Wilma’s action causes a state of affairs, *Will’s having poor health*—where health, or so Gardner indicates, is or involves a condition satisfying the requirement in clause (i)—such that, if Will existed and this state of affairs did not obtain, then he would have been better off with regard to the relevant condition.<sup>3</sup> It does not matter if Will would have never existed, and thus would not have had better health, if he

<sup>3</sup> Of course, this claim does not require Will to be intrinsically *badly* off with respect to his health.

did not have poor health: clause (ii) is satisfied anyway, as it only requires that Will would have been relevantly better off *if* he had existed without poor health. (The reference to the individual's existence in (ii)'s antecedent, I presume, is what gives the Existence Account its name.)

In my view, the Existence Account is a mistaken view of harming, and it works even less well as a solution to the non-identity problem. I shall make three interrelated critical remarks.

**(a) A Counterexample** The Existence Account yields, implausibly, that an action that causes something good for someone can harm him, even if the agent had no way of making things even better for him. Ironically, this point is illustrated by a case that Gardner puts forward as a counterexample to Elizabeth Harman's view (Harman 2004, 2009):

*Dim Vision.* Jones has been blind for many years as a result of a retinal damage. Recently, Dr. Smith has developed a surgical operation that can repair some but not all of the damage. Dr. Smith operates on Jones and improves his vision from a state of blindness to a state in which Jones can see, but not very well: Jones now has what we will call *dim vision*. (Gardner 2015a: 431)<sup>4</sup>

Obviously, Dr. Smith's operation does not harm Jones. According to Gardner, whereas Harman's view implies that it does (for reasons we need not go into), the Existence Account implies that it does not. Gardner claims that clause (ii) is not satisfied, for if "Jones were to exist without dim vision, then Dr. Smith would not have operated and Jones would be completely blind" (Gardner 2015a: 438).

But why should we assume with Gardner that Dr. Smith would not have operated if Jones had existed without dim vision? It seems no less reasonable to suppose that Dr. Smith would still have operated, but with a different—and perhaps even better—result with regard to Jones's vision. Presumably, Gardner is here relying on a "backtracking" view of causation, to the effect that A causes B only if A would not have taken place if it were not for B.<sup>5</sup> On this view, given that Dr. Smith's operation caused Jones to have dim vision, the operation would not have taken place if Jones did not have dim vision. Plausibly, it is then also true that if Jones had *existed* without dim vision—recall condition (ii)—then the operation would not have taken place. However, it still does not follow that if Jones had existed without dim vision, then Dr. Smith would not have operated on him; we are free to suppose that she would have performed an alternative, but similar operation instead.

Indeed, we can make the problem more striking by adding some details to the case, and modifying others. Suppose that Dr. Smith's operation causes blind Jones to have, not dim vision, but normal (good) vision. Suppose also that every time Dr. Smith operates, her assistant does some additional work on the patient's eyes—something that, due to the unreliability of the mechanisms involved, often makes no difference but sometimes improves the patient's vision even further. We can assume that Dr. Smith is determined to help Jones as much as possible, so that among the possible worlds in which Jones exists without normal vision, some world in which he has better than normal vision is nearer to the actual world than is any world in which he has worse than normal vision. Plausibly, the former, nearer world is one where Dr.

<sup>4</sup> *Dim Vision* is based on a similar case in Hanser (2009).

<sup>5</sup> See Gardner (2017b). I do not accept the backtracking view, but shall not argue against it here. See also Johansson & Risberg (2018).

Smith operates on Jones and the assistant's additional work happens to yield even better vision for Jones. (Again, to accommodate the above backtracking view of causation, we can suppose that in that world Dr. Smith performs an alternative but similar operation to the actual one.) In this variant of the case, Dr. Smith's action causes a state of affairs—*Jones's having normal vision*—such that, if Jones had existed and this state of affairs not obtained, then Jones would have been better off in some relevant respect.<sup>6</sup> So the Existence Account implies, implausibly, that Dr. Smith's operation, which takes Jones from no vision to normal vision, harms Jones.

A proponent of Gardner's position might try to deny that Dr. Smith's operation causes the state of affairs of *Jones's having normal vision*. But then what does cause it? (Surely something does.) One suggestion might be that *Jones's having normal vision* is caused by the assistant's (unsuccessful) work on his eyes. However, we can assume that the assistant's work does not satisfy the above "backtracking" condition: we can assume that it would have been performed even if Jones had existed without normal vision. Another suggestion might be that *Jones's having normal vision* is caused by whatever is causally responsible for the assistant's work being unsuccessful. However, there need be no such thing: we can assume that whether the additional work is successful or unsuccessful is undetermined, a matter of pure chance.

Some might think that it is actually a defensible implication that Dr. Smith's action harms Jones, assuming that Dr. Smith's alternative operation would have improved Jones's vision even more than her actual operation did. But we can simply stipulate that this assumption does not hold—and more generally, that none of Dr. Smith's alternative actions is such that, if it had been performed, then Jones would have had better vision (or been intrinsically better off in some other respect).<sup>7</sup>

If someone still insists that Dr. Smith's action harms Jones, that does not help much. For surely there is no moral reason against Dr. Smith's action. This suggests that the Existence Account, even if correct, would identify at most a morally innocent kind of harm, one that would not reveal any moral objection to the relevant action. But a morally innocent kind of harm is not what the harm-based approach to the non-identity problem promised to deliver (see Section 1).

**(b) The Harmfulness of Alternative Actions** A related criticism of Gardner's approach begins with the observation that in order for an action's harmfulness to provide a moral reason against it, it has to be the case that at least one of the agent's alternatives would not also be harmful (or at least not to the same or greater extent). A moral reason against an action *a* must speak in favor of performing some alternative action instead, and if *a*'s having a certain feature speaks

<sup>6</sup> Assuming, that is, that vision is or involves a condition with respect to which a person can be *intrinsically* better or worse off (as Gardner herself apparently assumes). If it does not, we can simply replace vision with something that more plausibly satisfies the requirement in (i), such as pleasure.

<sup>7</sup> Bontly (2016) defends a *contrastive* approach to causation and harm. Simplifying somewhat, Bontly's view of harm is that action *a* harms someone if and only if the agent's doing *a* rather than some appropriate contrast action *a*\* causes a state of affairs *s* rather than some appropriate contrast state of affairs *s*\*, and *s* makes the person worse off than *s*\* would have. This view fares no better than Gardner's in the modified vision case.

Admittedly, Dr. Smith's performing her actual operation rather than (say) refraining from operating at all does not cause *Jones's having normal vision* rather than *Jones's having better than normal vision* (as refraining from operating would not have given Jones better than normal vision). However, Dr. Smith's performing her actual operation rather than refraining from operating at all does cause *Jones's having normal vision* rather than *Jones's not having normal vision* (as the latter would obtain had she not operated at all, leaving Jones blind). And as before, we can suppose that the latter state of affairs would make Jones better off (giving him better than normal vision) than the former would. Hence, Bontly's view, too, entails that Dr. Smith's action harms Jones. Thanks to Ben Bradley for discussion here.

in favor of performing some alternative action instead, it cannot be that all alternatives to *a* have that feature as well. Now, consider our Wilma case. Without jeopardizing the intuition that Wilma's actual action is immoral, we can construct the case so that creating Liam, by the lights of the Existence Account, would also have been harmful. Of course, Liam would have been much happier than Will is; perhaps he would not have been intrinsically worse off than Will is in any respect. This, however, is compatible with the assumption that many states of affairs, which would be caused by creating Liam, would deprive Liam of enormous amounts of happiness. That is, each of a large number of states of affairs, which would be caused by creating Liam, would be such that if he had existed and it did not obtain, then he would have been much better off. Indeed, if we like, we can suppose that this holds to a greater extent in Liam's case than in Will's—that the difference between what Will actually gets and what he is deprived of is smaller than the corresponding difference in Liam's case. For instance, perhaps Liam would have a more impressive capacity to enjoy himself than Will does, so that relative to the nearest world in which Liam is created, possible worlds where he is enjoying himself *immensely* are fairly close. Surely these factors do not in any way reduce or even begin to counterbalance the moral objection to creating Will. Yet on these assumptions, the Existence Account yields that just as creating Will harms Will, creating Liam would harm Liam (and would even harm him *more*, on a natural way for advocates of the Existence Account to calculate degrees of harm). So again, even granting the Existence Account, and thereby the claim that Wilma's action harms Will, this provides no solution to the non-identity problem.

**(c) A Diagnosis** One diagnosis of these problems suggests itself. In fact, the diagnosis reveals a deeper problem behind not only Gardner's account, but any remotely similar version of the Will-focused view. Note first that the non-identity problem arises even on the supposition that Wilma's action does not cause anything intrinsically bad for Will. As Ben Bradley remarks (2012: 406), Wilma's action seems immoral even if we only assume that due to Will's disease, he is bound to receive only very minor intrinsic goods—for instance, because the disease gives him a severely limited capacity to enjoy life.<sup>8</sup> Presumably, then, on this construal of the case, if the states of affairs that Wilma's action causes are harms for Will, this is wholly because of the intrinsic *goods* that they *deprive* him of. Naturally, this is not itself objectionable; there is indeed such a thing as the harm of deprivation (death being the most prominent example). What is objectionable is the idea that *causing* such a state of affairs morally counts against an action, even in cases (such as the Wilma case and the vision case, in both their original and modified versions) where the relevant intrinsic goods are nothing that the agent *can bring about*—that is, in cases where none of the agent's available alternatives is such that, had it been performed, then the person would have received those intrinsic goods. We should all agree that an action's moral status at least partially depends on the *intrinsic* values for individuals of the states of affairs that it causes, and on the intrinsic values for individuals of the states of affairs that the action's respective alternatives would cause. However, we should deny that an action's moral status even partially depends, in addition, on whether those states of affairs *themselves* deprive anyone of goods (cf. Bradley 2009: 125). It does not morally count against an action that it brings about a scenario that in some respects is close to an even better scenario, at least if the latter scenario is one that the agent cannot bring about. Causing a near hit instead of a clear miss is not morally objectionable, at least unless one of the agent's available alternatives would have brought about a hit.

<sup>8</sup> Cf. the modest pleasures in the "Z" world figuring in Parfit's "repugnant conclusion" (Parfit 1984: 388).

### 3 Harming the Merely Possible Child: The Liam-Focused View

But the harm-based approach is not hopeless. There is a version of it that has been insufficiently explored—one that requires a controversial (though defensible) ontological view, but might be attractive to those who, unlike me, are prepared to accept that ontological view. However, before we get there (Section 4), it will be helpful to see what can be made of the simpler idea that Wilma’s action is wrong largely because it harms her counterfactual child, Liam (call this the “Liam-focused view”).

It may be thought that the Liam-focused view is ruled out already by our stipulation at the outset, that Wilma’s action has no significant effect on anyone distinct from her child. After all, that child is Will, not Liam. Here, however, it is important to note the “anyone” in the stipulation. Because Liam is never created, no one is Liam: it holds for anyone that he or she is not Liam. This consideration also suggests that although it may seem that the Liam-focused view is in conflict with Premise 1 (that Wilma’s action does not make anyone worse off than he or she would have otherwise been), it is rather a way of denying Premise 4 (that if Wilma’s action does not wrong anyone, then it is morally permissible). The idea is that although Wilma’s action does not (harm or) wrong anyone, it is still wrong, since it harms Liam. The idea, in other words, is that Liam is no one, but still harmed.

The Liam-focused view has one clear advantage over the Will-focused view. Plausibly, creating Will would have been permissible if the only alternative had been to create no one at all. This judgment seems particularly compelling in light of the fact (mentioned in Section 2) that we are free to assume, both in the original case and in this variant of it, that Wilma’s action does not cause anything that is *intrinsically* bad for Will, but merely gives Will a severely limited capacity to enjoy himself. Especially on this assumption, Wilma’s action seems morally permissible in the variant of the case (where the alternative is to create no one), and it still seems morally objectionable in the original case (where the alternative is to create Liam). Since Wilma’s action does the very same thing to Will in the variant of the case as in the original case, the permissibility of creating Will in the variant of the case suggests that the wrongness of creating Will in the original case is not primarily due to what it does to Will. Rather, it seems to be primarily due to the presence a *better alternative* (cf. Bradley 2012: 406).<sup>9</sup> By emphasizing Liam’s fate, the Liam-focused view is more in keeping with this consideration than is the Will-focused view.

But how can nonexistent Liam be harmed by Wilma’s action? The idea would likely be something like the following (cf. Holtug 2001: 375–77). For an existing person to be made worse off by an action, and hence harmed by it, it is sufficient that he or she never receives any intrinsic goods, but would have been well off if the action had not been performed. And harming someone in this way is often morally wrong. Now, Liam never receives any intrinsic goods, and would have been well off but for Wilma’s actual action. So why not say that Liam, too, is impermissibly harmed?

There is a good answer to this question, however (cf. Bykvist 2007: 343). It is not attractive to deny the following principle: in every case of harming, the relation *x harms y* obtains between the harming entity and the harmed entity. Nor is it attractive to deny the following principle: it is impossible to stand in any relation without (ever) existing. The Liam-focused

<sup>9</sup> Or perhaps: the presence of a better alternative that unlike the actually performed action avoids placing anyone at a low well-being level. We need not hold that it is morally objectionable to place someone at a high well-being level instead of placing someone at an even higher well-being level.

view, though, apparently forces us to reject at least one of these two principles. By contrast, no such rejection follows from the claim that an existing person can be harmed by being prevented from occupying a positive well-being level.

Any philosophical principle can be resisted with some degree of reasonableness, including these two. But a more promising response on behalf of the Liam-focused view would be that we have been too quick in denying Liam's existence. No doubt Liam is not an actual person. However, some philosophers hold that merely possible people exist—alternatively put, that something is a merely possible person. For example, on Timothy Williamson's version of this view (Williamson 2013), there is no possible world in which some entity is contingently existing.<sup>10</sup> Because there is a possible world where Liam exists, Liam exists in every possible world, including the actual one. The idea is thus not that Liam is an existing yet merely possible person in a sense that implies that whereas it is merely possible that Liam exists, Liam does exist. That idea can be falsified on the general grounds that if it is merely possible that  $p$  is true, then  $p$  is not true. Rather, on Williamson's view, Liam does exist, but is merely possibly a person, and thus merely possibly *someone*—and, indeed, merely possibly concrete. Aside from such modal properties, most of Liam's properties are extremely "thin": for example, *being an entity*, and *being numerically distinct from the number 2*. (In many other worlds, of course, Liam has lots of nonmodal "thick" properties, such as *being a person*.) Nonetheless, Liam is not nothing but something. Not someone, but something.

This ontological view is in some ways unorthodox and surprising, but Williamson shows it to have serious advantages over rival views, advantages that have nothing specifically to do with ethics. Appealing to it, then, is no ad hoc maneuver, or hopelessly unacceptable in any other way.

But although this ontological view ascribes existence to Liam, it remains the case that it denies Liam concreteness. While such denial is eminently sensible, it causes trouble for the Liam-focused view. For in order for Liam to be made worse off by Wilma's action, Liam's well-being level must be lower than it would have otherwise been. And that requires Liam to *have* a well-being level—at least a neutral one, a well-being level of zero. But here is an attractive principle (Herstein 2013): something has a well-being level, including a neutral one, only if it has some sort of capacity to have whatever is required for positive or negative well-being (e.g., pleasure or pain)—a capacity that involves more than merely having positive or negative well-being in some possible world in which the object is enormously different from the way it actually is. (The difference between concreteness and nonconcreteness certainly qualifies as enormous.) It is hard to see how a nonconcrete thing like Liam could have the requisite capacity.

Again, the relevant principle can be resisted. But once again, a better strategy would be to further refine the Liam-focused view. We'll do that next.

#### 4 Harming a Compound: The William-Focused View

An advocate of the harm-based approach to the non-identity problem might claim that Wilma's action harms, not Will or Liam, but some entity of which both of them are somehow constituents. (This works only if Liam exists—as nonexistent entities cannot be constituents of anything—and so let us continue to assume that Liam is a nonconcrete object.) Undoubtedly

<sup>10</sup> Williamson does not discuss the non-identity problem, and I have no reason to think that he would embrace the Liam-focused view.

some versions of this view would be nonstarters. One example is the suggestion that the *class* of Wilma's various possible children is harmed by her action (Peters 1999). Classes are abstract entities without well-being levels; indeed, unlike Liam, no class has a well-being level in any possible world at all, no matter how concrete its members. A preferable suggestion, it seems to me, appeals to a *mereological fusion* of Will and Liam, that is, an individual that has Will and Liam as its parts—an individual we might give the name “Will-Liam,” or simply “William.” William is thus *composed* of Will and Liam, in the very same way that Will, say, is composed of his upper half and lower half. (Unlike Will, William has a nonconcrete part, but the parthood relation is the same in both cases.) According to the “William-focused view,” then, Wilma's action is morally wrong largely because it harms William.<sup>11</sup> Although it is theoretically possible to grant the existence of each of Will and Liam while denying that there is something that they compose, an ontology generous enough to include Liam will likely have room for William as well.

William's having Will, who evidently has a well-being level, as a part makes it plausible that he himself has a well-being level.<sup>12</sup> (I take it that “he” is the right word, as William has the same body as Will.) Naturally, there are many properties that Will has and William lacks—for example, *being wholly concrete* and *being a proper part of William*. But it does seem that William and Will have the same *mental* properties. For since Will is a part of William, it appears that William must have the same brain, and consequently the same brain-states as Will (and the same physical surroundings, if that matters too). And whether or not it is metaphysically *necessary* that two things have the same mental features if they have the same brain-states (and physical surroundings), we should not deny that if two things *actually* have the same brain-states (and physical surroundings), then they *actually* have the same mental features. It seems, then, that William has whatever mental features that Will does. Someone might be tempted to protest here that William does not have any mental features at all, on the grounds that he has a part—Liam—that, unlike the brain, is not relevantly involved in the production of mental states. But that factor cannot prevent William from having mental states: after all, Will is not prevented from having mental states by having among his parts his hands, which, too, are not relevantly involved in the production of mental states. Others might want to protest that precisely because William has the same mental features as Will if William exists, we should deny that William exists. On the face of it, Will is the *only* thinker of his thoughts, and the only individual having his experiences. But while this protest might be correct as far as it goes, it is pointless to make it in this context. For I am not arguing that William exists—again, I do not embrace the William-focused view myself—but that he has the same mental states as Will *if* he exists.<sup>13</sup>

Now, if it is right that William, supposing he exists, has Will's psychology, then he would seem to have a well-being level as well. For surely having the same mental features as someone with a well-being level (in this case, Will) is sufficient for having a well-being level oneself. Indeed, William would seem to have the *same* well-being level as Will in the actual world. For instance, if hedonism is true, an individual's well-being level is wholly determined

<sup>11</sup> According to van der Zee and de Beaufort, the “collective of potential persons of parents A and B” can be harmed by the actions of A and B (2011: 452–53). The authors' talk of a “collective” is slightly unclear, but the William-focused view might be one way of developing their idea.

<sup>12</sup> Thanks to an anonymous reviewer for prompting me to clarify and elaborate on my reasoning (in the next few paragraphs) behind this claim.

<sup>13</sup> Some might think that this kind of implication is too implausible for the William-focused view to be worth investigating. But this is too strong a claim. See further (b) in Section V.

by his receipt of pleasure and pain, and William experiences whatever pleasure and pain that Will does. Similar remarks apply to most other theories of well-being, including many theories that deny that an individual's well-being depends solely on her mental states. Consider for example desire satisfactionism. William desires whatever Will desires, and so whatever satisfies or frustrates Will's desires, also satisfies or frustrates William's desires—and the other way around. Thus, although desire satisfactionists deny that an individual's well-being level is purely determined by her mental states, they should still say that William's well-being level is the same as Will's.

Analogous points hold for the nearest world where Wilma creates (i.e., makes concrete) Liam instead of Will. That is, just as William actually has Will's mental states, so William would have had Liam's mental states, and Liam's well-being level, had Liam been created.<sup>14</sup> Consequently, if Wilma had performed the alternative action, so that Will would be nonconcrete and Liam concrete, then William himself would have occupied a higher well-being level. Thus, Wilma's actual action makes William worse off than he himself would have otherwise been. And therefore, one might argue, even if Wilma's action harms neither Will nor Liam, it harms William, and impermissibly so.

Plausibly, being someone, rather than just something, is a psychological matter. In particular, being psychologically indiscernible from someone seems to be sufficient for being someone. Thus, if William is psychologically indiscernible from Will, who is someone, William, too, is someone—not just something, like Liam. Unlike the Liam-focused view, then, the William-focused view is most naturally regarded as a way of denying Premise 1 of the Booninian argument. That is, it is a way of claiming that Wilma's action does make someone worse off than he or she would have otherwise been.

It is instructive to compare the William-focused view with some other views in the same neighborhood. On one such view, even if Wilma's action does not harm any particular individual, it still harms "her child," understood *de dicto* (Hare 2007). For while it is actually the case that her child has a low well-being level, if she had created Liam it would have been the case that her child has a high well-being level. In contrast to this *de dicto* approach, the William-focused view appeals to harm *de re*—the harm done to a particular individual. On the William-focused view, it is true of someone (William) that he himself would have been better off had Wilma created Liam instead of Will.<sup>15</sup>

Another view, proposed by Parfit, appeals to the notion of a "general person." A general person, Parfit explains, is a group of possible people, one of whom will be created. According to Parfit, an action like Wilma's can harm a general person, such as the group of Will and Liam. But, he stresses, the phrase "general person" is "merely an abbreviation. Like *the Average American*, a general person is not a person" (2011, p. 220). Indeed, "[g]eneral people are not individuals" (2011, p. 236; emphasis removed). By contrast, William is a particular individual—and as indicated above, since Will is a person with whom William shares all mental features, there is little reason to deny him personhood. If Parfit is right, there is no individual (or person) *x* such that *x* is the Average American or a general person; by contrast, on the William-focused view, there is an individual (and person) *x* such that *x* is William.

<sup>14</sup> This presupposes that William has Liam as a part not only in the actual scenario, but in the relevant counterfactual scenario as well. While this may not be indisputable, it is plausible.

<sup>15</sup> I thus disagree with Boonin when he suggests (2014: 105n.) that views like the William-focused view reduce either to the impersonal approach or to the *de dicto* approach. None of these two approaches appeal to the harm done to any particular individual. For further comparison with the impersonal approach, see (c) in Section V.

Yet another view is that Wilma's action harms the group of all, or at least some, (actual) future people. For, it might be suggested, the very same group of future people that actually has Will as a member, would instead have had Liam as a member if Wilma had created Liam. This "group of future people" view, too, must be distinguished from the William-focused view. For instance, unlike proponents of the William-focused view, proponents of the group of future people view need not accept the existence of merely possible people (such as Liam), or any nonconcrete object at all. Conversely, proponents of the William-focused view need not accept the claim about group identity crucial to the group of future people view. Indeed, the William-focused view is perfectly compatible with the claim that any group of individuals essentially has exactly those members that it actually has—so that there is *no* possible world (let alone a nearby possible world) in which the group, which is actually the group of future people, contains Liam instead of Will.

### 5 Three Objections

I want now to consider three possible concerns about the William-focused view. As my aim is not to show that the William-focused view is correct, but only that this sort of view deserves more attention than it has received, my replies will be tentative.

**(a) Psychological, Physical, and Biographical Differences** If Wilma had created Liam, then Liam's (and hence William's) psychology would likely be very different from Will's. Liam's (and hence William's) body and biographical details would likely be very different as well. So while Wilma's actual action deprives William from occupying a higher well-being level, he would have had a very different life had she acted otherwise. Some might contend that this means that William is not, in the actual world, suitably related to the counterfactual scenario in which he is better off, so that Wilma's action does not harm William after all, or at least not morally impermissibly.

While that contention is not unreasonable, nor is its denial. After all, many would likely say that creating a perfectly ordinary person (with no nonconcrete part) can morally impermissibly harm that person if the agent's alternative action would have caused *that same person* a much better, albeit radically different life. Like many others, I prefer the life I actually have to one in which I am much happier but in which I am also in other respects dramatically different—and I have this preference not only for the sake of others, but also for my own sake. Maybe this attitude is even rationally permissible. But the lesson to draw from this might simply be that we can sometimes rationally permissibly prefer, even for our own sakes, an action that harms us to one that does not. This need not make the harmful action *morally* permissible.

**(b) Too Much Harm** Presumably, if there is such a thing as William, there are also many entities of which William is a part. For example, an ontology hospitable enough to include William should also include a compound of Wilma's *all* possible children: a compound containing everything that in some possible world is Wilma's child. This more comprehensive entity is composed of one concrete person, Will, and lots of nonconcrete objects in addition to Liam. If William has a well-being level, so does this more comprehensive entity. For just like William, this more comprehensive entity has the same brain-states as Will in the actual world, and therefore arguably the same psychology. Wilma's action thus makes this entity, too, worse

off than it would have otherwise been; for if she had acted otherwise, it would instead have had Liam's brain-states, and hence Liam's psychology and well-being level. And analogous remarks hold for other William-containing entities (such as a compound of Will, Liam, and just one of Wilma's other possible children; or a compound of Wilma's all possible children and one of Liam's possible children, etc.). It might be unappealing to say that such a large number of beings—this enormous number of entities having William as a part—are harmed by Wilma's action. We began by noting the difficulties in finding even one harmed individual; now we apparently have an excessive abundance.

Still, this kind of implication is already accepted by many metaphysicians with an ontologically generous mindset, even among those who do not countenance the existence of merely possible individuals. Friends of temporal parts, for example, think that if I—a temporally extended object—suffer at noon today, then my temporal parts I-at-noon-today, I-today, I-this-year, I-this-decade, and so forth, suffer as well. Many such philosophers would add that while this may sound strange, the worry can be mitigated by noting that in ordinary, nonphilosophical talk and thought, it would be perfectly appropriate to deny that I share my suffering at noon with an enormous number of individuals at my spatial location. This might be, for instance, because in such contexts it would be appropriate to count by *nonoverlapping* individuals (obviously, I overlap each of my temporal parts). If this kind of position is reasonable, it is not clear that it is unreasonable for advocates of the William-focused view to accept that, strictly speaking, Wilma's action harms not only William but also a vast number of individuals with whom he overlaps. In a nonphilosophical context, it might still be appropriate to deny that claim. Indeed, advocates of William-focused view could even say that in nonphilosophical contexts, it would be appropriate to deny that Wilma's action harms *anyone*, on the grounds that it does not harm any of the entities normally recognized in our ordinary, nonphilosophical discourse (such as Will). This would not jeopardize the William-focused view: harming an individual seems ethically significant whether or not that individual is normally recognized in our ordinary, nonphilosophical discourse.

Of course, many philosophers would insist that the sort of implication highlighted here—that there are strictly speaking many individuals harmed by Wilma's action, or sharing my suffering at noon—is a good reason not to be ontologically generous, never mind what advocates of the generous ontology allow us to say in ordinary, nonphilosophical contexts. But while I have sympathy for this reaction, it must be admitted that the reason in question is nowhere near decisive. The William-focused view, then, is in respectable company with regard to this sort of implication.

**(c) Impersonalist Implications** The William-focused view will likely share many moral implications with the impersonal approach. Specifically, the William-focused view apparently sacrifices one of the main supposed advantages of the harm-based approach—namely, its avoidance of the claim that we have a strong moral reason, or even a moral obligation, to create happy people even when we could instead create no one (Section 1). For it seems that if the William-focused view is correct, then for any actual person P who occupies a fairly low well-being level, there is a mereological fusion F of P and some merely possible person P\* who would be happy if created (made concrete). Assuming that P\*'s life would not somehow prevent P's life, if P\*

were created then F would consist of two concrete people, P and P\*. Since there would be a lot of happiness going on in P\*'s mind, it seems that creating P\* would have given F a higher well-being level than F actually has.<sup>16</sup> Now, if Wilma's action impermissibly harms William by placing him at a low well-being level instead of a much higher one, presumably we would also impermissibly harm F by not creating P\*. Again, it is often seen as a major problem for the impersonal approach to the non-identity problem that it has these sorts of implications, and a virtue of the harm-based approach that it does not. In a way, this advantage disappears on the William-focused view.<sup>17</sup>

But only in a way. For it must be pointed out that it seems considerably easier to swallow such an implication *if there really is* an individual who is made worse off by the action. After all, on the William-focused view, there is a living, breathing victim of flesh and blood (though one of its parts is utterly fleshless and bloodless) whom the action prevents from occupying a much higher well-being level. Unlike the impersonal approach, then, the William-focused view does not commit us to the claim that we are morally required to create a happy person even if we otherwise harm no one. A closely related point is that the William-focused view does not confer "moral value on states of affairs independently of persons," as David Heyd (1988: 154) accuses the impersonal approach of doing. Again, since Will is a person, there is little reason to deny that William is a person. He is not a merely possible person (and hence no person at all), but an actual person. For the same reason, moreover, unlike the impersonal view the William-focused view easily accommodates the much discussed "person-affecting restriction," according to which an outcome is better than another only if the former is better than the latter for someone.

I suggest, then, that if the kind of implication exemplified by the case of P and P\* is objectionable, then this is mainly because the underlying ontological view (assigning existence to P\* and F, and to Liam and William) is objectionable. To my mind, the intuitive unreality of merely possible people amounts to more than their mere failure to be concrete, and so (as I mentioned in Section 3) I myself am not prepared to accept that ontological view. However, it can hardly be said to be indefensible.

## 6 Conclusion

As I argued in Section 2, Gardner's approach is unpromising, and I gave grounds for doubting any remotely similar version of the Will-focused view. The Liam-focused view, the topic of Section 3, strikes me as even more unpromising. If we still want to identify a subject of harm in non-identity cases, it seems to me that the William-focused view deserves our attention. No doubt it is relatively drastic, and could be safely ignored if someone had already put forward some fairly unproblematic solution

<sup>16</sup> Are there any problems with this claim which arise from P and P\*'s having two distinct minds at one and the same time? If so, just assume that P's having mental states and P\*'s having mental states do not overlap in time. Perhaps there would then be no particular time at which F would be better off, had P\* been created, but it is hard to see why F would not have been better off *overall*.

<sup>17</sup> The same goes for the Liam-focused view, but as we have seen, that view has drawbacks more serious than this one.

to the non-identity problem (whether or not a version of the harm-based approach). But as we all know, that condition is far from being satisfied.<sup>18</sup>

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

## References

- Algander P (2013) Harm, benefit, and non-identity. Uppsala University, Doctoral dissertation
- Bontly T (2016) Causes, contrasts, and the non-identity problem. *Philos Stud* 173:1233–1251
- Boonin D (2014) *The non-identity problem and the ethics of future people*. Oxford UP, New York
- Bradley B (2009) *Well-being and death*. Oxford UP, New York
- Bradley B (2012) Doing away with harm. *Philos Phenomenol Res* 85:390–412
- Bykvist K (2007) The benefits of coming into existence. *Philos Stud* 135:332–362
- Gardner M (2015a) A harm-based solution to the non-identity problem. *Ergo* 17:427–444
- Gardner M (2015b) Review of Boonin (2014). *Notre dame philosophical reviews*, June 4, <<http://ndpr.nd.edu/news/the-non-identity-problem-and-the-ethics-of-future-people>>
- Gardner M (2015c) Well-being and the non-identity problem. In: Fletcher G (ed) *The Routledge handbook of philosophy of well-being*. Routledge, London, pp 429–438
- Gardner M (2016) Beneficence and procreation. *Philos Stud* 173:321–336
- Gardner M (2017a) On the strength of the reason against harming. *J Moral Philos* 14:73–87
- Gardner M (2017b) When good things happen to harmed people. *Ethical Theory and Moral Practice*, advance access. <https://doi.org/10.1007/s10677-017-9840-z>
- Hanser M (2009) Harming and procreating. In: Roberts M, Wasserman D (eds) *Harming future persons*. Springer, Dordrecht, pp 179–199
- 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
- Harman E (2004) Can we harm and benefit in creating? *Philos Perspect* 18:89–113
- Harman E (2009) Harming as causing harm. In: Roberts M, Wasserman D (eds) *Harming future persons*. Springer, Dordrecht, pp 137–154
- Herstein O (2013) Why ‘nonexistent people’ do not have zero wellbeing but no wellbeing at all. *J Appl Philos* 30:136–145
- Heyd D (1988) Procreation and value: can ethics deal with futurity problems? *Philosophia* 18:151–170
- Heyd D (1992) *Genethics: moral issues in the creation of people*. University of California Press, Berkeley
- Holtug N (2001) On the value of coming into existence. *J Ethics* 5:361–384
- Johansson J, Risberg O (2017) The preemption problem. *Philos Stud*. <https://doi.org/10.1007/s11098-017-1019-x>
- Johansson J, Risberg O (2018) The problem of justified harm: a reply to Gardner. *Ethic Theory Moral Prac*. <https://doi.org/10.1007/s10677-018-9912-8>
- Parfit D (1984) *Reasons and persons*. Oxford UP, Oxford
- Parfit D (2011) *On what matters: volume two*. Oxford UP, Oxford
- Peters PG (1999) Harming future persons: obligations to the children of reproductive technology. *South Calif Interdiscip Law J* 8:375–400
- Schwartz T (1978) Obligations to posterity. In: Sikora R, Barry B (eds) *Obligations to future generations*. Temple UP, Philadelphia, pp 3–13
- Shiffrin S (1999) Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory* 5:117–148
- van der Zee B, de Beaufort I (2011) Preconception care: a parenting protocol. A moral inquiry into the responsibility of future parents towards their future children. *Bioethics* 25:451–457
- Williamson T (2013) *Modal logic as metaphysics*. Oxford UP, Oxford

<sup>18</sup> For helpful comments I am grateful to three anonymous reviewers, and to participants at the "Harm: The Concept and Its Relevance" conference at Uppsala University in August 2016, especially Ben Bradley, Erik Carlson, Neil Feit, Molly Gardner, Matthew Hanser, Olle Risberg, and Simon Rosenqvist.