

ARTICLE

Asymmetry and Non-Identity

Per Algander^{1*} and Katharina Berndt Rasmussen^{2*}

¹Uppsala University and ²Institute for Futures Studies, Stockholm

*Corresponding author. Email per.algander@filosofi.uu.se
katharina.berndt.rasmussen@iffs.se

(Received 18 March 2018; revised 20 August 2018; accepted 20 December 2018)

Abstract

In this article we distinguish two versions of the non-identity problem: one involving positive well-being and one involving negative well-being. Intuitively, there seems to be a difference between the two versions of the problem. In the negative case it is clear that one ought to cause the better-off person to exist. However, it has recently been suggested that this is not so in the positive case. We argue that such an asymmetrical treatment of the two versions should be rejected and that this is evidence against views according to which it is permissible to cause the less well-off person to exist in the positive non-identity case.

Keywords: Non-identity problem; Asymmetry; Harm; Benefit; Population ethics

1. Two versions of the non-identity problem

A much-discussed case from Derek Parfit's *Reasons and Persons* is the non-identity problem.¹ In order to set up the problem, Parfit asks us to consider a case like the following:

The Non-Identity Case: A couple decides between having a child now or later. If they were to have a child now then this child, call her Ann, would have a life worth living. If they were to have a child later then they would have a different child, call her Beth, whose life would be well worth living. There are no other morally relevant considerations at stake.²

Intuitively, it seems clear that it is *better* if the couple waits.³ One problem, according to Parfit, is how to explain this intuitive judgement. In particular, Parfit argues that

¹Derek Parfit, *Reasons and Persons* (Oxford, 1984).

²We will assume that the Non-Identity Case is a two-outcome case, i.e. either the couple has a child now or later. Whether our arguments can be extended to non-identity cases where not having a child at all is an alternative is something we will leave for another occasion.

³This intuition has been taken for granted in much of the debate concerning the non-identity problem. See e.g. Julian Savulescu, 'Procreative Beneficence: Why We Should Select the Best Children', *Bioethics* 15 (2001), pp. 413–26; Gustaf Arrhenius, 'Can the Person-Affecting Intuition Solve the Problems in
© Cambridge University Press 2019. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

so-called person-affecting views in ethics cannot explain why it is better to wait since having a child now would not be worse for that child, and that such views should be rejected for this reason.

At first sight, it might also seem that the couple *ought* to wait. However, this idea has recently been challenged. David Boonin, for example, argues that common-sense morality implies that it is not wrong (i.e. permissible) for the couple to have a child now and that we should embrace this conclusion. Similar views have been expressed by Melinda Roberts and David Wasserman.⁴

These challenges are an interesting addition to the debate on the non-identity problem. In this article, we discuss some of the consequences of accepting that it is permissible for the couple to have a child now in the Non-Identity Case. In particular, we consider the following version of the case:

The Inverse Non-Identity Case: A couple decides between having a child now or later. If they were to have a child now then this child, call him Carl, would have a life worth *not* living. Carl's life would, overall, contain more of the things which make a life go badly than the things which make a life go well. If they were to have a child later then they would have a different child, call him Dave, whose life would also be worth not living. However, Dave's life would contain a better balance of good and bad things than Carl's life. It would contain fewer of the things which make a life go badly, but it would still, overall, contain more of the things which make life go badly than things which make a life go well. Again, there are no other morally relevant considerations at stake.

In this case it seems clear that it is *not* permissible for the couple to have a child now, that is, to make it the case that Carl exists.⁵ However, this version of the case is structurally identical to Parfit's original version of the case. To see the similarity, consider Tables 1 and 2, where the degree of a life being worth (not) living is represented by the number of '+' ('-'), while non-existence is represented by an empty box.

In either case, one person would exist if the couple decides to have a child now and another person would exist if they were to have a child later. Moreover, the child they would have were they to have a child now would be worse off than the child they would

Table 1. The Non-Identity Case

	Now	Later
Ann	+	
Beth		+++

Population Ethics?', *Harming Future Persons*, ed. M. A. Roberts and D. T. Wasserman (Dordrecht, 2009), pp. 289–314; Nils Holtug, *Persons, Interests, and Justice* (Oxford, 2010).

⁴David Boonin, *The Non-Identity Problem and the Ethics of Future People* (Oxford, 2014); Melinda Roberts, 'The Non-Identity Fallacy: Harm, Probability and Another Look at Parfit's Depletion Example', *Utilitas* 19 (2007), pp. 267–311; David T. Wasserman, 'Harms to Future People and Procreative Intentions', *Harming Future Persons*, ed. M. A. Roberts and D. T. Wasserman (Dordrecht, 2009), pp. 265–85.

⁵We do not take a stand on the question whether creating Dave is permissible or not, i.e. whether there might be genuine moral dilemmas. Our arguments do not turn on the moral status of the second alternative in either case.

Table 2. The Inverse Non-Identity Case

	Now	Later
Carl	---	
Dave		-

have had, had they decided to have a child later. In the light of this similarity, one might wonder *why* what the couple ought to do should differ in the two cases. Without a plausible account of a relevant difference between both cases, Boonin's view and its underlying common-sense morality are brought into question.

If there is no morally relevant difference between the two cases, then we should accept the following claim:

Symmetry: It is permissible for the couple to have a child now in the Non-Identity Case if and only if it is permissible for the couple to have a child now in the Inverse Non-Identity Case.

This claim is worth emphasizing since it will play an important role in the discussion below. Symmetry expresses the idea that the Non-Identity Case and its inverse counterpart should be treated in the same manner. However, note that it does not entail or express any kind of maximizing or utilitarian view of morality. It is, after all, compatible with Symmetry that it is permissible for the couple to have a child now in both versions of the non-identity case. What the claim rules out is a kind of 'mixed' view where non-identity cases involving lives worth living are treated differently from cases involving lives worth *not* living.

Our main aim in this article is neither to defend Symmetry, nor to argue against it. Rather, we are concerned with the viability of views implying that Symmetry is false: views according to which there is an important difference between the Non-Identity Case and the Inverse Non-Identity Case. Our main conclusion is that such views have such counterintuitive implications that avoiding Symmetry seems very costly.

II. Simple axiological asymmetry

A straightforward attempt to explain the difference between the two cases is based on a distinction between a person's well-being (here: the degree to which her life is, or is not, worth living) and the impersonal value of an outcome. While the value of an outcome plausibly depends in some way on the distribution of well-being in that outcome, there are several ways to spell out this dependence.

The simplest way of connecting well-being and the impersonal value of an outcome is to say that the contributive value of well-being always *equals* the amount of well-being.⁶ On this view, creating Ann and Carl is worse than creating Beth and Dave, respectively, since the former exist at lower well-being levels than the latter. This view – call it Simple Axiological Symmetry – thus fails to account for an axiological difference between both cases.

⁶The contributive value of a life is the difference in impersonal value which that life makes to an outcome. On classical utilitarianism the contributive value of a life always equals the well-being in that life. On other views, e.g. prioritarianism, the contributive value of a life can be greater, or smaller, than the well-being in that life.

In order to account for the purported axiological difference, the following view might seem more promising:

Simple Axiological Asymmetry: Adding people at positive well-being levels does not make an outcome better, but adding people at negative well-being levels makes an outcome proportionally worse.

On this view, the positive well-being of Ann and Beth does not contribute to the value of the outcome. Assuming that the two outcomes are comparable, it follows that the outcome where Ann exists with positive well-being and the outcome where Beth exists with more positive well-being are equally good, other things being equal.⁷ In the inverse version, however, that Carl exists with negative well-being and that Dave exists with less negative well-being are not equally good, because negative well-being makes an outcome proportionally worse. By disregarding the contributive value of positive well-being, but not that of negative well-being, this view gives us an axiological difference between the two cases.

This view does not contradict Symmetry by itself, since Symmetry is a claim about what is permissible, while Simple Axiological Asymmetry is a claim about the value of outcomes. To contradict Symmetry, we can add the following deontic principle:

Simple Deontic Principle: An action is permissible if and only if its outcome is at least as good as any of its alternatives' outcomes, other things being equal.

Together, the Simple Deontic Principle and the Simple Axiological Asymmetry entail that it is permissible for the couple to have a child now in the Non-Identity Case, and impermissible for the couple to have a child now in the inverse version of the case, thereby contradicting Symmetry.

However, Simple Axiological Asymmetry is not very plausible. John Broome has offered an argument which can be used against the first conjunct of Simple Axiological Asymmetry.⁸ Consider the following three outcomes:

- A: No one exists.
- B: Eve exists with a life worth living.
- C: Eve exists with a life well worth living.

Because Eve enjoys positive well-being in both *B* and *C*, it follows from Simple Axiological Asymmetry that *A* is equally as good as *B* and that *A* is equally as good as *C*. These claims entail that *B* is equally as good as *C*.⁹ A plausible principle when ranking outcomes where the same people exist is the following:

⁷That two outcomes, O_1 and O_2 , are comparable means here that one of the three standard value relations holds: O_1 is better, worse or equally as good as O_2 . This assumption can be questioned and we will consider views which do not rely on it below.

⁸John Broome, *Weighing Lives* (Oxford, 2004), p. 147.

⁹We here assume that 'is equally as good as' is transitive: if *A* is equally as good as *B*, and *B* is equally as good as *C*, then *A* is equally as good as *C*.

Pareto: If an outcome O_1 is better (worse) than O_2 for someone, and at least as good for everyone else, then O_1 is better (worse) than O_2 , other things being equal.

It follows from Pareto that C is better than B (because it is better for Eve and worse for no one). Hence Simple Axiological Asymmetry is incompatible with Pareto. Of these two, Pareto seems to be the more plausible one. We should therefore reject Simple Axiological Asymmetry.

Note that Broome's argument can also be used against other versions of Simple Axiological Asymmetry. Consider for example the following satisficing version of the view:

Satisficing Axiological Asymmetry: Positive well-being below well-being level w makes an outcome proportionally better, but positive well-being above w does not make an outcome better. Negative well-being makes an outcome proportionally worse.

On this view, negative well-being always makes an outcome worse and positive well-being below a certain threshold level makes an outcome proportionally better. However, further additions of positive well-being above the threshold level do not make an outcome better. To see that this view is inconsistent with Pareto we need only suppose that Eve's well-being in both B and C is above the threshold. If it is, then it follows from Satisficing Axiological Asymmetry that B and C are equally good (again, assuming that the two outcomes are comparable).

It might be objected that this argument begs the question against any version of axiological asymmetry. Satisficing Axiological Asymmetry is after all consistent with Pareto below the threshold, and that additions of well-being above the threshold do not make an outcome better is simply what the view states. Objecting that the view is inconsistent with Pareto above the threshold might therefore look question-begging.

This is not a very strong objection. As mentioned at the beginning of this section, it is very plausible that the value of an outcome depends, at least in part, on the distribution of well-being in that outcome. Pareto and Satisficing Axiological Asymmetry are two ways to spell out this dependency. As we have seen, these two cannot both be true. The question then is which view is most plausible as a view about the relation between a distribution of well-being and the value of an outcome. Here it seems to us that Pareto should be preferred. To claim that one view is more plausible than another for a specific purpose is not to beg the question against that other view.

A more promising view is to note that it does not follow from Simple Axiological Asymmetry alone that adding people with positive well-being results in an equally good outcome; this also hinges on the assumption that the two outcomes are comparable. Simple Axiological Asymmetry just states that adding people at positive well-being levels does not make an outcome better. This is compatible with two distributions of positive well-being being *incomparable*.¹⁰

¹⁰This view seems to be suggested by Jan Österberg, 'Value and Existence: The Problem of Future Generations', *Odds and Ends*, ed. S. Lindström, R. Sliwinski and J. Österberg (Uppsala, 1996), pp. 94–107. See also Wlodek Rabinowicz, 'Broome and the Intuition of Neutrality', *Philosophical Issues* 19 (2009), pp. 389–411.

That two valuable things are incomparable is often defined by saying that it is neither the case that one is better than the other, nor that they are equally good.¹¹ Simple Axiological Asymmetry, understood in terms of incomparability, then amounts to the view that making it the case that Ann exists and making it the case that Beth exists are incomparable in value, other things being equal.

Typical examples of incomparability are when things exemplify radically different values, such as Sartre's example of the student facing a choice between fighting for his country or taking care of his mother. A problem for the incomparability view is that the two non-identity cases do not seem to involve radically different values. Furthermore, even if it is granted that causing Ann to exist is incomparable in value to causing Beth to exist, it still needs to be explained why this is *not* the case with respect to Carl and Dave. If the two outcomes in the Non-Identity Case are incomparable then one would expect the outcomes in the inverse version to be incomparable too. If the latter is denied, we need some explanation of this difference between the cases. If none is given, the reply is simply an attempt to explain the axiological asymmetry with an unexplained asymmetry in the comparability of outcomes. This is not satisfactory.

A related view is to say that it is indeterminate, or vague, whether it is better to cause Ann to exist than to cause Beth to exist.¹² However, this view faces the same problem as the incomparability view. If it is indeterminate in the Non-Identity Case whether it is better to cause Ann or Beth to exist, then we need some reason to think that it is not indeterminate in the inverse version whether it is better to cause Carl or Dave to exist. Simply insisting that it is indeterminate in the positive case, but not in the negative case, is unpersuasive.

The indeterminacy view faces further problems when offered as a defence of Simple Axiological Asymmetry. Consider how the indeterminacy of betterness works in other areas. It seems plausible that if it is indeterminate whether x is better than y , then it is possible that there is a version of x , x^+ , which is determinately better than y , and a version of x , x^- , which is determinately worse than y . However, if the indeterminacy of betterness is to explain Simple Axiological Asymmetry, it cannot work in this way. Assume that it is indeterminate whether causing Beth to exist (x), is better than causing Ann to exist (y), in the Non-Identity Case. Then there is a version of x , with Beth's life *not* worth living, which is determinately worse than y . But there is no version of x , with Beth's life even more worth living, which is determinately better than y . This difference between how the indeterminacy of betterness works in other areas and how it is supposed to work in the non-identity case suggests that the indeterminacy of betterness alone will not help explain the difference between the Non-Identity Case and its inverse counterpart.

It may be suggested that we could still avoid Symmetry with the help of a 'capped' incomparability or indeterminacy view.¹³ On these views, there is a range of well-being levels such that adding a person at one of these levels makes the outcomes incomparable

¹¹The terminology surrounding (in)comparability is rather diverse. Broome, for example, uses the term 'incommensurate' for what we call 'incomparability' (Broome, *Weighing Lives*, p. 22).

¹²Versions of this view are suggested by Broome, *Weighing Lives*, and by Campbell Brown, 'Better Never to Have Been Believed: Benatar on the Harm of Existence', *Economics and Philosophy* 27 (2011), pp. 45–52. See Rabinowicz, 'Neutrality', for a critical assessment of Broome's view.

¹³For such a 'critical level' indeterminacy/incomparability view, see e.g. Broome, *Weighing Lives*, and Rabinowicz, 'Neutrality'. We owe this suggestion to an anonymous reviewer of this journal.

or indeterminate. Symmetry will therefore only hold for non-identity cases where the well-being of the people involved is above or below the range. We think that such views are problematic for the following reasons. First, regardless of where the range of indeterminacy or incomparability is set, we can restate the Non-Identity Case with well-being levels above the range (leaving the inverse case unchanged). Second, capped views run counter to the intuitive appeal of treating the Non-Identity Case differently from the inverse version. On these views, Symmetry would still hold for some *very high* and *very low* levels of well-being. But, that the two cases should be treated differently only seems *more* compelling at higher levels of well-being.

Thus, without further qualifications, Simple Axiological Asymmetry does not provide a satisfactory account for a relevant difference between the two versions of the non-identity case. By mapping outcome value quite simplistically onto positive and negative well-being levels, respectively, this view seems to take into account the only difference there is between the cases. However, there might be other features, which become recognizable only within a less simplistic framework.

A noteworthy feature of non-identity cases is that they involve people with certain modal properties. Could some modal restriction on counting well-being explain the difference between our two cases?¹⁴ It should be clear that none could. Consider that Ann and Beth are both contingent people, since neither of them will exist in both outcomes.¹⁵ However, this is also true of Carl and Dave. Since the cases are structurally identical regarding the modal properties, modal restrictions cannot establish a relevant difference.

There are, however, other ways to bring out differences between the two cases by considering more complex patterns of their respective features. The following section examines one such suggestion, which will turn out to be a dead end. But as such, it makes for an informative detour, pointing us in a more promising direction.

III. The person-affecting restriction

According to the person-affecting approach, the value of an outcome depends on the *comparative value* of each outcome *for each person*.¹⁶ More specifically, this approach satisfies the following restriction:

The Person-Affecting Restriction: An outcome O_1 is better (worse) than an outcome O_2 only if O_1 is better (worse) than O_2 for someone.

Can this restriction, together with the Simple Deontic View, be used to avoid Symmetry?

¹⁴Peter Singer, *Practical Ethics* (Cambridge, 1993), p. 103, suggests that only people who will exist independently of an act matter to the evaluation of the act. On this view, the well-being of people not 'necessary in the circumstances' are fully discounted. For different actualist views, see Frank Jackson and Robert Pargetter, 'Oughts, Options, and Actualism', *The Philosophical Review* 95 (1986), pp. 233–55; John Bigelow and Robert Pargetter, 'Morality, Potential Persons and Abortion', *American Philosophical Quarterly* 25 (1988), pp. 173–81; Josh Parsons, 'Axiological Actualism', *Australasian Journal of Philosophy* 20 (2002), pp. 137–47.

¹⁵Everyone is obviously contingent in the sense that they could have not existed. 'Contingent' is here understood as 'contingent relative to the available alternatives'.

¹⁶For a defence of the view see Melinda Roberts, 'Is the Person-Affecting Intuition Paradoxical?', *Theory and Decision* 55 (2002), pp. 1–44, and Holtug, *Persons*, ch. 5.

In order to address this question, the restriction itself needs to be clarified. First, can an outcome O_1 be better (or worse) for a person than O_2 if the person exists in only one of these outcomes? If it cannot, then the Person-Affecting Restriction rules out the possibility that it is better (or worse) to create Carl in the Inverse Non-Identity Case. In order for this view to get off the ground, it will have to be assumed that an outcome O_1 can be worse than O_2 for a person even if the person does not exist in both outcomes, regardless of whether O_1 or O_2 were to obtain.¹⁷

Second, in the Inverse Non-Identity Case, the Person-Affecting Restriction does *not* imply that having a child later (Dave) is better than having a child now (Carl). This is because the restriction only states a necessary condition for when an outcome is better (worse) than another. In order to derive the conclusion that it is better to have a child later in the Inverse Non-Identity Case, the restriction needs to be strengthened.

One plausible strategy is to combine it with the Pareto-principle mentioned above. Unfortunately, adding Pareto is not enough. Having a child later is, admittedly, better for the person who would exist if the couple were to have a child now (Carl). But it is not at least as good for everyone else, since it is *worse* for the person who would exist if the couple were to wait and have a child later (Dave).

To save this proposal, it might be suggested that, in addition to Pareto, the 'better for' and 'worse for' relations can only be instantiated if all their relata exist. If the couple were to have a child now, then that outcome would be worse for Carl because he would exist and that outcome would be worse for him than the outcome where he doesn't exist. On the other hand, were the couple to have a child now then this outcome would *not* be better for Dave because Dave does not exist in that outcome. Therefore, if the couple were to have a child now in the Inverse Non-Identity Case then this would be worse for Carl and better for no one, hence it would be worse than having a child later.¹⁸

However, the same argument can be made if the couple chooses to have a child later. Having a child later would be worse for Dave because he would exist, and this outcome would be worse for him than the outcome where he does not exist. That outcome would *not* be better for Carl because Carl does not exist in that outcome. If the couple were to have a child later then this would be worse for Dave and better for no one, hence it would be worse than having a child now. This proposal therefore implies that whatever the couple does, the outcome would be *worse* than the alternative.

Could this proposal account for a relevant difference between the two cases? Regarding the choice whether to cause Ann or Beth to exist, this view implies that if the couple were to cause Ann to exist then this would be better for Ann and worse for no one, hence it would be better. The same argument can be made if the couple chooses to cause Beth to exist. Causing Beth to exist would be better for Beth and worse for no one. It would therefore be better to cause Beth to exist. This view therefore implies that whatever the couple does, the outcome would be *better* than the alternative. The view, combined with the Simple Deontic Principle, therefore implies that Symmetry is false. According to Symmetry, it is permissible to cause Ann to exist if

¹⁷We grant this controversial assumption for the sake of the argument. For recent defences of the assumption, see Jens Johansson, 'Being and Betterness', *Utilitas* 22 (2010), pp. 285–302; Gustaf Arrhenius and Wlodek Rabinowicz, 'Better to Be Than Not to Be', *The Benefit of Broad Horizons*, ed. H. Joas and B. Klein (Leiden, 2010), pp. 65–85.

¹⁸There are many problems with this reply. Some of the difficulties which are not raised here are discussed by Arrhenius and Rabinowicz, 'Better to Be'.

and only if it is permissible to cause Carl to exist. The person-affecting view discussed here implies that the couple will do something permissible, whatever they do, in the Non-Identity Case, and something impermissible, whatever they do, in the Inverse Non-Identity Case.

While the qualified Person-Affecting Restriction avoids Symmetry, it does so at a significant cost. Suppose that the couple decides to cause Dave to exist in the Inverse Non-Identity Case. Causing Dave to exist is then worse than causing Carl to exist. According to the Simple Deontic Principle it is therefore impermissible to cause Dave to exist. However, note that if it is worse to cause Dave to exist, then it is better to cause Carl to exist. This view therefore implies, absurdly, that *if the couple were to cause Dave to exist, then it is permissible to cause Carl to exist*. This conclusion is very implausible. Causing Carl to exist should be impermissible, regardless of what the couple does.¹⁹ The Person-Affecting Restriction will therefore not provide an adequate solution to our puzzle.

In being partly successful, i.e. managing to establish an asymmetry between our two cases, this view points to a different approach, which might do better. Note that the problem with the Person-Affecting Restriction stems from the Simple Deontic Principle. What the Person-Affecting Restriction seems to get right is that by causing Carl to exist we *harm* him, and this has moral relevance. A remaining puzzle is whether we *benefit* Ann (or Beth) by creating her, and whether that has any moral relevance.

In the next two sections we will consider an account framed in terms of individual comparisons of well-being: harms and benefits. We will first clarify what we mean by 'harm' and 'benefit' (section IV), and then consider whether there is a deontic principle, formulated in terms of harm and benefit, which can account for the difference between the two non-identity cases (section V).

IV. Harm and benefit

The underlying idea which we will explore in this section and the next is that what explains the difference between the two versions of the non-identity case has to do with the moral relevance of harming and benefiting people. Bringing Carl into existence harms him, and does so to a greater degree than bringing Dave into existence would have harmed Dave. Similarly, Ann is benefited by being brought into existence, and Beth would have benefited to a greater degree had she been brought into existence instead. However, it does not follow merely from these remarks that Beth or Dave ought to be brought into existence.

Making claims about harm and benefit in non-identity cases is notoriously controversial.²⁰ The problem is that harm and benefit seem to be comparative notions: in order for an event to harm (benefit) a person, the event must make the person worse (better) off. Furthermore, it seems plausible that the relevant comparison is

¹⁹Moreover, this version of the person-affecting restriction makes the normative status of an action depend on whether it is performed (assuming the Simple Deontic Principle). It is highly controversial whether such violations of 'Normative Invariance' are acceptable. See Erik Carlson, *Consequentialism Reconsidered* (Dordrecht, 1995); Krister Bykvist, 'Violations of Normative Invariance: Some Thoughts on Shifty Oughts', *Theoria* 73 (2007), pp. 264–83; Krister Bykvist, 'Prudence for Changing Selves', *Utilitas* 18 (2006), pp. 264–83.

²⁰Derek Parfit famously argued that the non-identity problem, as he conceived it, cannot be solved by appealing to person-affecting notions such as harm (Parfit, *Reasons and Persons*, pp. 357–66).

with what would, or could, otherwise have been the case. One way to capture these intuitive ideas is the following condition:

The Counterfactual Condition: An event e harms (benefits) a person p only if it would have been better (worse) for p had e not occurred.²¹

Applied to non-identity cases, the Counterfactual Condition raises the question – again – whether it can be better, or worse, for a person to exist than never to have existed. A quick argument that such comparisons do not make sense goes as follows. A person would have been better (worse) off had an event e not occurred if and only if the person would have had more (less) well-being had e not occurred. However, a person would not have had any amount of well-being (not even zero) had she never existed. Therefore, never having existed cannot be better or worse for a person than existing with either positive or negative well-being.

The upshot of this quick argument is that a person cannot be benefited or harmed by being caused to exist. This is relevant to our present purpose since it shows that any attempt to reject Symmetry by appealing to harm and benefit will have to either (i) claim that existence can be better or worse for a person than never having existed, or (ii) reject the Counterfactual Condition.²²

However, the quick argument might be too quick. Melinda Roberts and Nils Holtug, among others, have argued that we should ascribe zero (or neutral) well-being to people in outcomes where they do not exist.²³ On their view, existence would be better (worse) for a person than non-existence just in case the person would have positive (negative) well-being, were she to exist. In the Non-Identity Case, causing Ann (or Beth) to exist would therefore benefit her. Likewise, causing Carl (or Dave) to exist would harm him.

A troublesome consequence of Roberts's and Holtug's view is that it allows for non-existent people to be benefited and harmed. To illustrate, the view implies not only that Ann would be benefited by being caused to exist, but also that she would be *harmed* if she were not caused to exist. This is troublesome because there is a sound metaphysical reason for denying that claim. If a person does not exist in a possible world w , then she cannot exemplify any properties or stand in any relations in w , and being benefited or harmed by an event is to exemplify a property. This problem can be dealt with by adding an 'existence requirement' to the Counterfactual Condition:

Existence Requirement: An event e harms (benefits) a person p only if p exists and e occurs.

²¹The Counterfactual Condition is sometimes formulated in terms of 'could' rather than 'would'. The 'could' formulation is preferable when dealing with cases with more than two alternatives, or when it is indeterminate what would have been the case had e not occurred. Since we will consider simple cases with only two alternatives, there should not be a difference between the two formulations.

²²Boonin, for example, would have to go for (i) since he accepts the Counterfactual Condition (Boonin, *The Non-Identity Problem*, ch. 1). In the remainder of this section we will discuss attempts to reject Symmetry based on the Counterfactual Condition. In final section (VI) we will consider views which reject it.

²³Melinda Roberts, 'The Asymmetry: A Solution', *Theoria* 77 (2011), pp. 333–67; Holtug, *Persons*, pp. 129–30. An alternative approach to defending (i), which does not rely on ascribing well-being to non-existent people, is offered by Johansson, 'Being and Betterness' and Arrhenius and Rabinowicz, 'Better to Be'.

The Counterfactual Condition and the Existence Requirement provide two necessary conditions for harm and benefit. In what follows we will assume that they are also jointly sufficient. With this additional assumption, it follows that causing a person to exist at a positive well-being level benefits that person, while causing a person to exist at a negative well-being level harms that person (assuming that the person would not have existed otherwise).

With the addition of the Existence Requirement we avoid the implication that failing to cause a person to exist, who would have existed at a negative (positive) well-being level had she existed, is to benefit (harm) that person. This means that the problem of omission, which usually haunts counterfactual analyses of harm, is avoided in non-identity cases. The problem is due to the idea that harm and benefit are just ‘two sides of the same coin’, such that failing to benefit someone is equivalent to harming her.²⁴ The Existence Requirement undercuts this idea in non-identity cases. Consider the Inverse Non-Identity Case. Here, the view under discussion implies that we would harm Carl if we were to cause him to exist. If harms and benefits are merely two sides of the same coin, we would have to say that we then benefit Carl by causing him not to exist. The Existence Requirement saves us from this strange conclusion.

It should be noted that it is still an open question what the implications for permissibility are. The following section explores how, and at what cost, Symmetry – the view that it is permissible to cause Ann to exist if and only if it is permissible to cause Carl to exist – could be avoided.

V. Harming, wrongdoing and impermissibility

An easy way to achieve a deontic asymmetry between the ordinary and inverse non-identity case is to say that harms are relevant for permissibility, while benefits are not. According to the counterfactual analysis of harm described in the previous section, it follows that Carl and Dave would be harmed in the Inverse Non-Identity Case, were they to be caused to exist, and that neither Ann nor Beth would be harmed in the Non-Identity Case, were they to be caused to exist. This difference in harming is, on the view under consideration, what explains the difference at the deontic level between the two cases.

A view along these lines is suggested by Boonin, who considers the following two principles:

- P4: If an act does not harm anyone, then the act does not wrong anyone.
 P5: If an act does not wrong anyone, then the act is permissible.²⁵

Together with a counterfactual analysis of harm,²⁶ these two principles entail that it is permissible to cause Ann to exist and that it is permissible to cause Beth to exist.

²⁴A view along these lines is suggested by Roberts, ‘Asymmetry’. Seana Shiffrin discusses the view (without endorsing it), describing harm and benefit as ‘two ends of a scale’ (Seana Shiffrin, ‘Wrongful Life, Procreative Responsibility, and the Significance of Harm’, *Legal Theory* 5 (1999), pp. 117–48, at 121).

²⁵Boonin, *The Non-Identity Problem*, p. 6. Boonin uses ‘not morally wrong’ instead of ‘permissible’. We adapt his terminology to our framework since this does not affect the content of Boonin’s arguments. It should also be noted that while Boonin defends P4 and P5 (chs 5 and 6), a weaker and more plausible version of these principles would include an ‘other-things-being-equal’ clause.

²⁶Boonin does not explicitly endorse the counterfactual analysis as we formulate it here. He merely assumes that making people worse off than they would otherwise have been is *necessary* for harming,

However, they do not entail that it is impermissible to cause Carl (or Dave, for that matter) to exist in the Inverse Non-Identity Case. The view, therefore, does not contradict Symmetry. In order to yield the desired result in the inverse case – that it is impermissible to cause Carl to exist – it has to be assumed that harming is sufficient for impermissibility. A problem for such a version of Boonin's view is that it also implies that causing Dave to exist is impermissible. This then commits us to the view that there are genuine moral dilemmas, as instantiated by the Inverse Non-Identity Case: the couple will end up harming someone, and therefore doing something impermissible, whatever they do.

Accepting the possibility of moral dilemmas may seem a small price to pay, but there is a deeper problem here. Impermissibility does not allow for degrees, and thus creating Carl and creating Dave seem to be on a par, deontically speaking. This is counter-intuitive. There is an important difference between these two alternatives: surely, creating Carl, with a life well worth not living, is morally less acceptable than creating Dave, with a life just barely worth not living. This version of Boonin's view cannot account for this difference.

To avoid this problem, we can revise Boonin's view further in the following way:

Wronging: If an act φ harms a person p (to some degree d) then φ wrongs p (to some degree d^*).

Wrongness: An act φ is impermissible if and only if there is an alternative act ψ such that φ wrongs those who would exist if φ were to be performed more than ψ would wrong those who would exist if ψ were to be performed, other things being equal.

This view – let's call it *Complex Asymmetry* – makes impermissibility sensitive to degrees of wronging, where the degree to which a person is wronged is a function of the degree to which she is harmed. Complex Asymmetry can explain why it is permissible to cause both Ann and Beth to exist, and why it is impermissible to cause Carl to exist – without committing us to moral dilemmas or failing to account for the moral difference between causing Carl to exist and causing Dave to exist. In short, it provides a principled and intuitively plausible way to deny Symmetry.

Complex Asymmetry has certain similarities with the Simple Axiological Asymmetry considered above. Both views are ways of capturing the idea that the goods that could befall Ann and Beth lack moral relevance while the bads which could befall Carl and Dave are morally relevant.

However, Complex Asymmetry seems superior to the Simple Axiological Asymmetry since it is consistent with Pareto-like considerations. To illustrate, Complex Asymmetry implies that if one alternative is worse for someone, and at least as bad for everyone else, then it is impermissible. This means that in *same-people cases* – cases where the same people will exist whatever one does – Complex Asymmetry implies that there is no morally relevant difference between cases where the alternatives contain lives worth not living and cases where they contain lives (well) worth living. Consider the following same-people versions of the two non-identity cases. A couple decides between providing their child, call her Ann-Beth, with either (A) limited resources, giving her a life worth living, or (B) plenty of resources, making her life

not that it is sufficient. However, the stronger version of the counterfactual analysis is required in order to avoid Symmetry.

well worth living. Another couple decides between providing their child, let's call him Carl-Dave, with (C) limited resources, giving him a life well worth not living, or (D) plenty of resources, making his life just barely worth not living. In both cases, Complex Asymmetry implies that, all else being equal, it is impermissible to choose the alternative containing less well-being, be it positive or negative. It is only in *non-identity cases* that Complex Asymmetry implies that there is a morally relevant difference between causing people with lives worth not living to exist and causing people with lives worth living to exist.²⁷

Note that until now we have only considered simple non-identity cases, in a double sense: *pure non-identity cases*, containing alternatives with contingent people only, with the *same number* of individuals (only one). This raises the question whether Complex Asymmetry has equally acceptable conclusions in *mixed* (possible and actual people) and in *different number* cases. As it turns out, these cases highlight that on this view, benefits matter too little, and harms matter too much. Consider the following:

The Non-Identity Different Number Case (Table 3): A policymaker decides between policy A and policy B. A will cause one person to exist, with a life worth living. B will cause a billion (different) people to exist, each with a life well worth living.

Table 3. The Non-Identity Different-Number Case

	Policy A	Policy B
Person 1	+	
Person 2		+++
...		...
Person 1 000 000 001		+++

Complex Asymmetry implies that, just as in the case with Ann and Beth, it is permissible to adopt policy A, rather than policy B. This is clearly counterintuitive. In fact, this different number case gives us an *extreme super-repugnant conclusion* for this view.²⁸ The view's core feature of making only harms, but not benefits, relevant for impermissibility repugnantly permits choosing to cause only one person with low (positive) well-being to exist over causing a billion people with high well-being to exist.

A defender of Complex Asymmetry may now want to claim that this is just a version of the Non-Identity Case. Since the view is an attempt to accommodate the intuition that it is permissible in that case to cause the person with less well-being to exist, it

²⁷David Boonin suggested in his comments to this article that perhaps we should say that a person is wronged if the person's life is worth not living (personal communication, July 2018). This would allow us to say in the Inverse Non-Identity Case that creating either Carl or Dave would wrong that person. However, we think that our formulation above is overall preferable to Boonin's suggestion. For one thing, this suggestion does not account for degrees of wrongdoing, as discussed previously. Moreover, it is unclear whether the view suggested by Boonin has the above-mentioned advantages concerning Pareto and same-people cases. For example, it does not explain why we would wrong Ann-Beth by creating her with less well-being than she could have had, while Complex Asymmetry can accommodate this intuition.

²⁸This is an extreme deontic version of Nils Holtug's (axiological) 'Super-repugnant Conclusion' (Holtug, *Persons*, p. 254).

may seem question-begging to complain that the view implies that it is permissible to cause the person with less well-being to exist in the present case.²⁹

However, this reply is misguided. The intuition that it is impermissible to choose policy A in the present case is on an equal footing with our intuition regarding the Non-Identity Case. If our intuitions about the Non-Identity Case support an asymmetric view, such as Complex Asymmetry (by *inference to the best explanation*), then our intuitions about the Non-Identity Different Number Case undermine such a view equally. Moreover, there is much more at stake in the Non-Identity Different Number Case, concerning vast numbers of possible people, compared to individual reproductive – small-scale – choices. The pressure to get it right is arguably higher in large-scale cases, and hence the counterintuitive result should worry defenders of Complex Asymmetry.

A possible reply to this objection is to restrict Complex Asymmetry to same-number cases. The challenge, recall, was to find a view which could account for our intuitions in two particular same-number cases: the Non-Identity Case and the Inverse Non-Identity Case. A defender of Complex Asymmetry could argue that her view should only be applied to such same-number cases, and that the view would have to be further developed, or complemented, in order to deal with different number cases. Still, without any systematic proposal of how such a development could be achieved, this reply remains unsatisfactory.

Moreover, this view also has severely counterintuitive consequences in same-number cases. Consider the following:

The Mixed Non-Identity Case: A woman undergoing fertility treatment has three fertilized eggs planted into her uterus. It then turns out that the fertilized eggs have a genetic condition, making them sensitive to the intake of a certain drug shortly before implanting. If the woman takes a dose of the drug immediately, one fertilized egg will implant and develop to be an extremely well-off person, One; another egg will implant but develop an impairment that will make the resulting person Two's life slightly worth not living, while the third egg will be so badly damaged that it will not even implant. If, on the other hand, the woman refrains from taking the dose immediately, the first egg will implant and develop a slight impairment, making the resulting person's life quite well worth living, the second egg will not implant, while the third will implant and develop to be an extremely well-off person.

According to Complex Asymmetry, the benefit to Three of being brought into existence does not matter for permissibility. The harm to Two of being brought into existence does matter, but it is outweighed by the harm that One would suffer, were Two caused to not exist through the choice of not taking the drug. According to Complex Asymmetry it is therefore permissible to take the drug, and impermissible to not take the drug, since the degree to which One would be wronged, were the drug not taken, would be less than the degree to which Two would be wronged, were the drug to be taken. What Complex Asymmetry amounts to in this case is therefore that the only permissible alternative in this situation is to take the drug, making it the case

²⁹Boonin, *The Non-Identity Problem*, pp. 119–20, discusses a case like this. He raises some doubts whether his view actually has the implication that it is not impermissible to choose policy A, but also defends the view that it is not obviously impermissible to choose policy A on similar grounds to the one we raise here.

that one extremely well-off child (One) and a somewhat badly off child (Two) will exist, when one could have had an extremely well-off child (Three) and a quite well-off one (One). This is clearly counterintuitive.

The problem which this case illustrates is that Complex Asymmetry fails to take into account benefits to people who would not otherwise have existed. The benefit to One, if the woman chooses to take the drug, outweighs the harm to Two. However, the harm to One, if the woman decides to not take the drug, counterintuitively outweighs the *much greater* benefit to Three.³⁰

David Boonin has suggested that the Mixed Non-Identity case can be solved by appealing to whether an act makes things worse for a set of individuals, rather than individuals considered separately.³¹ On this view, if an act makes those who would exist, were the act to be performed, not worse off (worse off) than this set of individuals would be if the act were not performed, then the act is permissible (wrong). However, while appealing to sets of individuals (and thus comparing *aggregates* of well-being) may make the view more plausible in some cases, it does not make Boonin's view any more plausible with respect to our Mixed Non-Identity Case. It is still the case that the set consisting of One and Two would not be worse off, were the drug to be taken, than the set of them would be, were the drug not taken. To see this, suppose that the symbols in Table 4 represent units of well-being. Taking the drug gives the set of One and Two four units (five '+' minus one '-'), while not taking the drug gives the set consisting of only One (due to Two's non-existence) only three units. Taking the drug is thus permissible on this view – just as it is according to Complex Asymmetry.

Table 4. The Mixed Non-Identity Case

	Taking the drug	Not taking the drug
One	+++++	+++
Two	-	
Three		+++++

Of course, Boonin's view also entails that not taking the drug is permissible: it gives the set of One and Three eight units, whereas taking the drug gives the set of them only five units. Thus, taking the drug is not the *only* permissible alternative. However, as with Complex Asymmetry, it is the implication that it is permissible to take the drug which is counterintuitive. That it would also be permissible to not take the drug does not mitigate this strange result.

Could we make this an easier bullet to bite by considering how the agent might justify her choice to Two? One idea is that if the agent took the drug then she could justify

³⁰The Mixed Non-Identity Case also presents a problem for views which are similar to Complex Asymmetry, e.g. the view defended by Shiffrin. Shiffrin's view implies that it is permissible for the woman in the Mixed Non-Identity Case to take the drug, thereby making it the case that she has one child who is extremely well off (One) and one child who is slightly badly off (Two). On her view, the benefit to One matters morally since by taking the drug the woman 'removes or prevents' One from being much worse off. If the woman were not to take the drug then this would harm One and benefit Three (Shiffrin, 'Wrongful Life', p. 124). However, the benefit to Three should be disregarded since it does not remove or prevent any harm. The overall verdict, on Shiffrin's view, therefore seems to favour taking the drug.

³¹For such a view, see Boonin, *The Non-Identity Problem*, ch. 6 and appendix F, where he discusses what he calls the 'exclusive version of the moderate principle'.

the bad state which Two is in by saying: 'It was permissible for me to take the drug because even though this made you a little worse off than not existing, it made your sibling, One, much better off than they would otherwise have been.'³² Such a justification seems very implausible to us. First, Two (if sufficiently well-informed) could plausibly respond that had the agent not taken the drug, thereby making One worse off, then the agent could have offered One an even better justification than the agent is offering to Two: 'It was permissible for me not to take the drug because even though this made you, One, a little worse off, it gave your sibling, Three, a life extremely well worth living, and it spared another potential sibling a life worth not living.' This response illuminates why taking the drug is not permissible. Second, the agent's original justification seems far too permissive. What it amounts to, in effect, is a view where it is permissible for an agent to perform an action as long as the balance of counterfactual benefits and harms is positive, even if there is an alternative with an even better balance of benefits and harms. This view therefore suffers from the same problem as Complex Asymmetry in, for example, the Non-Identity Different Number Case.

To sum up: Complex Asymmetry gives us the desired asymmetric result in pure, same-number (small-scale) non-identity cases. However, the view has seriously counterintuitive implications in different-number cases. Furthermore, it will not help to restrict the view to same-number cases since, as the Mixed Non-Identity Case illustrates, the view has severely counterintuitive implications in those cases as well.

Our suggested diagnosis is that Complex Asymmetry fails – just as Boonin's alternative suggestion – because it gives certain benefits too little weight. We have thus reason to believe that a view according to which benefits lack moral significance should be rejected. The problem, in short, is that such a view is committed to counterintuitive, indeed repugnant, conclusions.

In the next section, we will consider a final attempt to avoid Symmetry. The view we will explore is similar to Complex Asymmetry in that it is harm-based, but it differs in that it rejects the Counterfactual Condition.

VI. Rejecting the counterfactual condition

We have already noted that a significant assumption in the discussion above is the Counterfactual Condition. According to this condition, an act harms a person only if the person would have been better off had the act not been performed. In this section we will briefly consider whether Symmetry can more plausibly be denied if we reject this condition on harms and benefits.³³ An example of one such view is the following:

The Intrinsic Causal View: An event harms (benefits) a person *p* if and only if it causes *p* to be intrinsically badly off (well off).

Views along these lines have been suggested by several authors.³⁴ The Intrinsic Causal View has certain theoretical advantages over the Counterfactual Condition

³²This move was suggested by Boonin (personal communication, July 2018).

³³Note that by rejecting the Counterfactual Condition one does not necessarily commit to a 'non-comparative' view of harm. An anonymous reviewer of this article, for example, suggested the view that an event harms a person if and only if the event's occurrence is worse for the person than its non-occurrence. This view would plausibly be a comparative, but not a counterfactual, view of harm.

³⁴See e.g. Elizabeth Harman, 'Can We Harm and Benefit in Creating?', *Philosophical Perspectives* 18 (2004), pp. 89–113; Krister Bykvist, 'The Benefits of Coming Into Existence', *Philosophical Studies* 135

regarding the explanation it offers for how causing a person to exist can benefit or harm that person. On the Counterfactual View, causing someone to exist can only harm that person if that person would have been better off had she not been caused to exist. However, as mentioned previously, it is doubtful whether this comparison – how well off the person is in the actual world compared to how well off she would have been had she never existed – even makes sense.

This worry is avoided on the Intrinsic Causal View. According to this view, what matters to harming and benefiting is whether an act causes a person to be intrinsically badly or well off, and this can certainly be the case in procreation.

Applied to the non-identity cases, this view implies that a couple would harm Carl by causing him to exist and that the couple would harm Dave by causing him to exist. However, the view also implies that a couple would benefit Ann by causing her to exist and that they would benefit Beth by causing her to exist. On the level of benefits and harms this view does not imply any asymmetry between the Non-Identity Case and the Inverse Non-Identity Case.

In order to defend an asymmetric treatment of the two cases one would therefore have to claim that the moral relevance of harms differs in some important way from the moral relevance of benefits. One way to do this would be to appeal, again, to Complex Asymmetry:

Wronging: If an act φ harms a person p (to some degree d) then φ wrongs p (to some degree d^*).

Wrongness: An act φ is impermissible if and only if there is an alternative act ψ such that φ wrongs those who would exist if φ were to be performed more than ψ would wrong those who would exist if ψ were to be performed, other things being equal.

With these normative views, the Intrinsic Causal View avoids the problem with the Mixed Non-Identity Problem. However, this view has other counterintuitive consequences. Consider the following:

The Same-Person Case (Table 5): A woman is expecting a child, One. If she takes the drug while she is pregnant, then One will develop to be a well-off person. If she does not take the drug, then One will develop an impairment which will have a severe negative effect on One's overall well-being. However, One's life will still be worth living with this impairment.

Table 5. The Same-Person Case

	Taking the drug	Not taking the drug
One	+++++	+

In this case, the Intrinsic Causal View implies that the woman would not harm One if she chooses not to take the drug. It follows from Complex Asymmetry that the woman does not wrong One and that she hence does nothing wrong if she chooses

(2007), pp. 335–62; Derek Parfit, 'Future People, the Non-Identity Problem, and Person-Affecting Principles', *Philosophy and Public Affairs* 45 (2017), pp. 118–57.

not to take the drug. This is clearly absurd. By not taking the drug the woman is doing something morally wrong.

There are several ways to avoid this absurd result while sticking with the Intrinsic Causal View. One such way is to claim that benefits have *some* moral significance, though perhaps not equal to the significance of harming.³⁵ We could then account for the intuition that the woman acts wrongly by not taking the drug in the Same-Person Case by appealing to the greater benefit to the child of taking the drug.

However, if benefits matter morally – although less than harms – then there is no relevant difference between our initial two cases, the original and the Inverse Non-Identity Case. If the fact that Beth would benefit more than Ann in the ordinary case is of moral importance then, other things being equal, it is difficult to see how it could be permissible to cause Ann to exist. This view therefore implies Symmetry.

VII. Concluding remarks

Some philosophers have argued that it is permissible to cause either person to exist in the Non-Identity Case. However, no one, we believe, holds that it is permissible to cause either person to exist in the Inverse Non-Identity Case. If one holds these two beliefs then one rejects Symmetry: it is permissible for the couple to have a child now in the Non-Identity Case if and only if it is permissible for the couple to have a child now in the Inverse Non-Identity Case. In this article we have discussed a number of ways in which one could defend this asymmetric view. The conclusions to be drawn from this discussion are the following.

First, we should reject the view that there is an *axiological* asymmetry which could ground the difference between the two Non-Identity Cases. In particular, we have shown that person-affecting axiologies are not sufficient to warrant the rejection of Symmetry. A noteworthy upshot of this conclusion is that value-based theories of morality should accept Symmetry.

Second, it is possible to reject Symmetry by appealing to the different moral significance of harms and benefits. We outlined one such view, Complex Asymmetry, inspired by recent work by David Boonin. We showed that such a view has severely counter-intuitive consequences both in different number cases and in same number cases. Whether Complex Asymmetry is more plausible than Symmetry is therefore doubtful.³⁶

Author ORCIDs.  Per Algander, 0000-0002-2333-106X; Katharina Berndt Rasmussen, 0000-0002-6805-4908

³⁵This view has been suggested by Harman, 'Can We Harm'.

³⁶The authors would like to thank Jens Johansson, Krister Bykvist, Gustaf Arrhenius and the attendees at the Stockholm June Workshop in Philosophy 2017 for their many helpful comments on drafts of this article. Thanks also to attendees at the research seminar in practical philosophy at Uppsala University and the philosophy seminar at the Institute for Futures Studies in Stockholm, where early drafts of this article have been presented. Berndt Rasmussen gratefully acknowledges the hospitality of the ANU School of Philosophy, where parts of this article were written in 2017. This article was written as a part of the research project Harm: The Concept and its Relevance, generously funded by the Bank of Sweden Tercentenary Foundation (dnr: P14-0212:1).