

Asymmetry and Non-Identity

Per Algander

1 Two versions of the Non-Identity Problem

A much discussed case from Derek Parfit's *Reasons and Persons* (1984) is "the Non-Identity Problem". In order to set up the problem, Parfit asks us to consider a case like the following:

The Non-Identity Case: A couple could decide to have a child now or later. If they were to have a child now then this child, call her Ann, would have a life barely 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.

The problem, according to Parfit, is to explain why the couple ought to wait. The problem is not *whether* it is permissible or not for the couple to have a child now or later, but *why* it is permissible for the couple to have a child later, and impermissible for them to have a child now. In particular, Parfit argues that so-called person-affecting views in ethics cannot explain why it is impermissible for the couple to have a child now since having a child now would not be worse for that child.¹

Parfit's assumption that it is impermissible for the couple to have a child now has recently been questioned. David Boonin (2014), for example, argues that common-sense morality implies that it is permissible for the couple to have a child now and that we should embrace this conclusion.

¹ Roberts (2007) argues that it seems to be assumed in Parfit's argument against person-affecting views that it is a two-outcome case. That is, either the couple has a child now or later. Not having a child at all is assumed to not be an alternative. In what follows I will only consider two-outcome versions of non-identity cases. Whether my arguments can be extended to non-identity cases where not having a child at all is an alternative is something I will leave for another occasion.

Boonin's view is an interesting and refreshing addition to the debate regarding non-identity cases and the Non-Identity Problem. In this paper I will discuss some of the consequences of accepting that it is, at least sometimes, permissible for the couple to have a child now in non-identity cases. In particular, I will consider the following version of the Non-Identity Case:

The Inverse Non-Identity Case: A couple could decide to have 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 makes 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.

In the Inverse Non-Identity Case it seems clear that it is not permissible for the couple to have a child now, thereby making it the case that Carl exists. However, this version of the case is structurally identical to Parfit's original version of the case. In both cases, 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, and the child they would have were they to have a child now would be worse off than the child they would have had, had they decided to have a child later. In the light of this similarity it is surprising that what the couple ought to do should differ in the two cases. This difference therefore requires an explanation.

There are two ways in which one could reply to this call for an explanation. First, it could be argued that there is an axiological

asymmetry between the two cases. That is, there is a difference in the *impersonal value* between making it the case that a person with a good life exists and making it the case that a person with a bad life exists. This view will be spelled out in further detail below, but in short the idea is that we are not making the world any better by causing a person with a good life to exist, though we *are* making the world worse by causing a person with a bad life to exist.

The second possible view is that there is a deontological asymmetry between the two cases. On this view, the difference between the two cases has to do with a difference in the relevant moral norms. On this view, the relevant moral norms are such that they prescribe that it is permissible for the couple to have a child now in the Non-Identity Case but not in the inverse version.

Regarding the first view, it is clear that it would have to be supplemented by some “bridge”-principle taking us from values to oughts and permissibility. After all, the problem is about what the couple ought, or are permissible, to do. This does not make the first view parasitic or dependent on the second however. An attractive feature of the first view is that it purports to explain the difference between the two cases while keeping the first-order normative principles intact. Exactly what these principles are, however, is not something I will discuss in this paper.

If both these views fail then one should claim that there is no difference between the two cases.² That is, we should accept the following claim

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

² Examples of authors who have argued for views which amount to a reply of this kind include Parfit (1984), Savulescu (2001), Belshaw (2003) and Bradley (2013).

This claim is worth emphasising 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 maximising 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 positive well-being are treated differently than cases involving negative well-being. In particular, proponents of the view that there is the kind of asymmetry in our duties to future people will probably be less inclined to accept Symmetry.³

My main aim in this paper is not to defend Symmetry. I will be concerned with the viability of views which imply that Symmetry is false: views according to which there is an important difference between the Non-Identity Case and the Inverse Non-Identity Case. In fact, Symmetry is controversial because it commits one to either a very demanding or a very permissive view of morality when it comes to procreation. The claim implies that either a couple ought to have the most well-off child they can, or that it is permissible for them to have any child; even a child who would suffer excruciating pain throughout her whole life, were she to exist. Neither of these view fits very well with our considered judgements about what is and is not permissible when it comes to procreation.

2 Simple axiological asymmetry

A straightforward attempt to explain the difference between the two cases is based on a distinction between a person’s level of well-being and the 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. Most

³ See for example McMahan (2009) and Roberts (2011).

importantly, we need not assume that how much a given amount of well-being contributes to the value of an outcome always equals the amount of well-being.

One way to develop this idea is the following axiological asymmetry:

Simple Axiological Asymmetry: Positive well-being does not make an outcome better, but negative well-being makes an outcome proportionally worse.

On this view, the positive well-being of Ann and Beth in the original Non-Identity Case do not contribute to the value of the outcome. That Ann exists with positive well-being and that Beth exists with more positive well-being are therefore two equally good outcomes, other things being equal. In the inverse version, on the other hand, that Carl exists with negative well-being and that Dave exists with negative well-being are not equally good, because negative well-being makes an outcome proportionally worse.

John Broome (2004, p. 147) has offered the following argument against the first conjunct of Simple Axiological Asymmetry. Consider the following three outcomes: A , no one exists; B , Eve exists with a life worth living; and C , Eve exists with a life well worth living. Because Eve enjoys positive well-being in both B and C , and positive well-being does not make an outcome better, 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 . However, a plausible principle when ranking outcomes where the same people exist is the following:

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 .

It follows from Pareto that C is better than B (because it is better for Eve), and hence Simple Axiological Asymmetry is incompatible

with this principle. Of these two, Pareto seems to be the more plausible one and we should therefore reject Simple Axiological Asymmetry.

Note that Broome's argument also rules out other potential 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, increases in 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 does not make an outcome better. To see that this view is also inconsistent with Pareto we need only suppose that Eve's well-being in B and C is above the threshold. If it is, then it follows from Satisficing Axiological Asymmetry that B and C are equally good.

One objection to Broome's argument is that it begs the question against any version of axiological asymmetry. In particular, Satisficing Axiological Asymmetry is after all consistent with Pareto below the threshold, and that further additions of well-being above the threshold does not make an outcome better is simply what the view is. Objecting that the view is inconsistent with Pareto above the threshold might therefore look question-begging.

I do not think this is a very strong objection. As mentioned at the beginning of this section, it is very plausible that the value of an outcome depends in some way on the distribution of well-being in that outcome. Pareto and Satisficing Axiological Asymmetry are two ways to spell out this dependency. However, we have seen that 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 me, Pareto ought to 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 reply to Broome's objection is to note that it doesn't follow from the Simple Axiological Asymmetry that two outcomes with different distributions of positive well-being are equally good. Simple Axiological Asymmetry is the view that positive well-being does not make an outcome better, but that is compatible with two distributions of positive well-being being *incomparable*.⁴

That two valuable things are incomparable is often defined by saying that it is not the case that one is better than the other and it is not the case that they are equally good.⁵ The 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 Ann exists are incomparable in value, other things being equal.

Typical examples of incomparability are things which exemplify radically different values, such as Sartre's example of the student who faces 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 be cases where radically different values are being compared. Comparing one person's well-being to another person's well-being does not seem to be the kind of case where we should expect incomparability. 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. That is, if the Non-Identity Case is a case of incomparable outcomes, then we need some

⁴ This view seems to be suggested by Österberg (1996). See also Rabinowicz (2009).

⁵ The terminology surrounding (in)comparability is rather diverse. Broome (2004, p. 22), for example, uses the term "incommensurate" for what I call "incomparability".

explanation of why the Inverse Non-Identity Case is not.

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. However, this view faces the same problems as the incomparability view. If it is indeterminate in the Non-Identity Case whether it is better to cause Ann or Beth to exist, why is not indeterminate in the Inverse Non-Identity Case whether it is better to cause Carl or Dave to exist?⁶

The indeterminacy view faces further problems when offered as a defence of Simple Axiological Asymmetry. Consider how vagueness works for other relations such as “is more bald than”. If it is vague whether someone, x , is more bald than y then it seems that we can imagine versions of x , x^+ , which are determinately more bald than x and versions of x , x^- , which are determinately less bald than x . The indeterminacy of betterness, however, cannot work in this way if the vagueness of “better for” is to explain Simple Axiological Asymmetry. On the vagueness view, we cannot make an outcome determinately better by making it the case that a person with positive well-being exists, but we can make an outcome determinately worse by making it the case that a person with negative well-being exists. This difference between how indeterminacy works in other areas and how it is supposed to work for betterness suggests that the indeterminacy of betterness will not help explain the difference between the Non-Identity Case and its inverse counterpart.

Simple Axiological Asymmetry will not be able to explain the difference between the two version of the Non-Identity Case. We therefore need to look elsewhere for an explanation.

⁶ A version of this view is suggested by Broome (2004) and Brown (2011). See Rabinowicz (2009) for a critical assessment of Broome’s view.

3 Modal partiality and the person-affecting restriction

A noteworthy feature of the Non-Identity Case is that it involves people with certain modal properties. Ann and Beth are both possible people since neither of them will exist whatever the couple does. It might therefore be suggested that only actual people, people who are “necessary in the circumstances”, or some other modal restriction can explain the difference between the two cases.⁷

In general, views of this kind can be seen as putting forward some modal property M and arguing for the following principle:

Modal Partiality: The value of an outcome is fully determined by the well-being of people with modal property M .

A common problem for all these views is that whether M is exemplified by a person does not track the person’s well-being. They all therefore fail to explain the difference between the two Non-Identity Cases. For example, the property *being a person who will exist independently of whether an act ϕ is performed* is not exemplified by anyone in either the Non-Identity Case nor in the inverse version. Since the well-being of people with this property do not count toward the value of an outcome, according to this view, choosing to have a child now or later leads to equally good outcomes in both versions of the problem.

Analogous arguments can be formulated for actualism. Suppose that there is a version of actualism which implies that the couple would not be making the world better by having a child later in the Non-Identity Case. But, whether this claim is true cannot depend on the fact that the Non-Identity Case involves positive well-being.

⁷ Singer (1993, p. 103) suggests that only people who will exist independently of an act matter to the evaluation the act. For different actualist views, see Jackson & Pargetter (1986), Bigelow & Pargetter (1988) and Parsons (2002).

If actualism has this consequence in this case, then it will have the parallel consequence in the Inverse Non-Identity Case as well.

An alternative, but similar, approach is to suggest that the difference between the two Non-Identity Cases has to do with person-affecting consideration:⁸

The Person-Affecting Restriction: If an outcome O_1 is better (worse) than an outcome O_2 then O_1 is better (worse) than O_2 for someone.

In order for this view to have any hope of explaining why it is worse to have a child now in the Inverse Non-Identity Case we have to assume that an outcome O_1 can be worse than O_2 for a person even if the person does not exist in both outcomes. Otherwise, the restriction would imply that the alternative in the Inverse Non-Identity Case are equally good.⁹

The Person-Affecting Restriction has some desirable consequences with respect to the Inverse Non-Identity Case. For example, it follows from the restriction that having a child now (cause Carl to exist) cannot be better than having a child later (cause Dave to exist). However, it does not follow that having a child later is better than having a child now. 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 strengthening of the Person-Affecting Restriction is to add the Pareto-principle mentioned above. Unfortunately, adding Pareto is not enough for it to follow that it is better to have a child later rather than now in the Inverse Non-Identity Case. Having a

⁸ The term “person-affecting” appears already in Parfit (1984). For a defence of the view see Roberts (2003).

⁹ Though controversial, I will grant this assumption in this context for the sake of argument. For a recent defence of this assumption, see Johansson (2010) and Arrhenius & Rabinowicz (2010).

child later is 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; it is *worse* for the person who would exist if the couple were to wait and have a child later (Dave).

It might be objected that the better for-relation can only be instantiated if all its relata exists. If the couple chooses to have a child now then this is worse for Carl because he exists and this outcome is worse for him than the outcome where he doesn't exist. However, it is not better for Dave because Dave does not exist. If the couple chooses to have a child now in the inverse Non-Identity Case then this is worse for Carl and better for no one, hence it is worse than having a child later.

There are many problems with this reply.¹⁰ For present purposes it is enough to note that this view does not imply that there is a difference between the two cases. Regarding the choice whether to cause Ann or Beth to exist it follows from these principles that if the couple chooses to cause Ann to exist then this is better for Ann and worse for no one. Hence it is better to cause Ann to exist. This view therefore implies Symmetry.

The view just which was just rejected highlights the possibility of combining Modal Partiality and the Person-Affecting Restriction in an attempt to explain the difference between the two versions of the Non-Identity Case. A schematic formulation of this class of views would look like this:

The Combined View: If an outcome O_1 is better (worse) than an outcome O_2 then O_1 is better (worse) than O_2 for some person p with modal property M .

It should be clear by now that the Combined View fails for the very same reasons as the previous two views. However M is specified, we have no reason to think that it will track any difference between

¹⁰ Some of the difficulties which are not raised here are discussed by Arrhenius & Rabinowicz (2010).

positive and negative well-being. That is, if M is specified such that it is worse to have a child now rather than wait in the Inverse Non-Identity Case, then the Combined View will also imply that it is worse to have a child now in the ordinary Non-Identity Case.

In general we can conclude that Modal Partiality and the Person-Affecting Restriction face the very same problem. In order for it to be worse to have a child now rather than later in the Inverse Non-Identity Case, but not in the ordinary version of the case, there has to be some feature of the inverse version which is not present in the ordinary version which can account for this difference. However, none of the views we have considered in this section identifies such a feature. Furthermore, because the cases are so similar, it is unclear what this feature could be. We should therefore conclude that an explanation of the difference between the two versions of the Non-Identity Case in axiological terms is not forthcoming.

4 Harm and benefit

The failure of axiological approaches to explaining the difference between the two versions of the Non-Identity Case is perhaps not very surprising. An alternative approach, which I will discuss in the following sections, is that we should explain the difference between the two cases in terms of a normative rather than an axiological difference. On the axiological approach some kind of bridge principles linking value to norms had to be taken for granted. The approach which will be discussed in the following sections is that the difference between the two cases is normative rather than axiological. More precisely, the hypothesis is that the difference between the two versions of the Non-Identity Case can be accounted for by considering the normative relevance of harms and benefits.

Making claims about harm and benefit in non-identity cases is notoriously controversial.¹¹ The problem is that harm and benefit

¹¹ Parfit (1984, pp. 357-366) famously argued that the Non-Identity Problem, as he con-

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 with what would, or could, otherwise have been the case. One way to capture this intuitive idea is the Counterfactual 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.¹²

The question which is raised by the Counterfactual Condition is whether it can be better, or worse, for a person to exist than to never 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.¹³

However, the quick argument might be too quick. It has been

ceived it, cannot be solved by appealing to person-affecting notions such as harm. See also Feinberg (1986)

¹² The 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. I will only consider simple cases with only two alternatives however, so there should not be a difference between the two formulations.

¹³ Boonin, for example, would have to go for (i) since he accepts the Counterfactual Condition. See Boonin (2014, ch. 1).

argued by Roberts (2011) and Holtug (2010), among others,¹⁴ that we should ascribe a well-being of zero to people who do not exist. On their view, existence would be better (worse) for a person than not existing just in case the person would have positive (negative) well-being, were she to exist. Causing p to exist with positive well-being would therefore benefit p , since p would then have positive well-being. Likewise, causing a person to exist with negative well-being would be to harm that person.

A troublesome consequence of Roberts' and Holtug's view is that it allows for non-existent people to be benefited and harmed. 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 a person p only if p exists and e occurs.

The Counterfactual Condition and the Existence Requirement are two necessary conditions for harm and benefit. In what follows I will assume that they are also jointly sufficient.¹⁵ With this additional assumption, it follows that causing a person with positive well-being to exist benefits that person, and causing a person with negative well-being to exist harms that person. However, it does not follow that *failing* to cause a person to exist, who would have had positive well-being had she existed, is to harm that person.

¹⁴ An alternative approach to defending (i), which does not rely on ascribing well-being to non-existent people, is offered by Johansson (2010) and Arrhenius & Rabinowicz (2010).

¹⁵ It is not necessary to assume a counterfactual analysis of harm in order to argue that we can harm a person by causing her to exist. For example, according to what we might call the *intrinsic causal view* an act (or event) harms a person iff the act (event) causes something intrinsically bad for the person. On this view, causing someone to exist with a life worth not living is plausibly to cause something intrinsically bad to happen to that person, and therefore causing her to exist harms her. Adopting the causal view will not make a difference to the discussion below however.

It should be emphasised that one is not endorsing Symmetry merely by accepting the analysis of harm and benefit just sketched. According to this view, the Non-Identity Case involves benefits and the Inverse Non-Identity Case involves harms, but it is an open question what the moral significance of these benefits and harms are. In order to reject Symmetry we would have to find some feature of the benefits in the Non-Identity Case which can serve as a ground for discounting their moral significance. However, this feature must not be present in inverse version.

Several possible views of the moral relevance of harms and benefits can be dismissed outright. For example, the view that harms matter more than benefits does not support a rejection of Symmetry. Since benefits still matter morally—although less than harms—on this view, there is no difference between the two non-identity cases. If the fact that Beth would benefit more than Ann in the ordinary non-identity problem is of moral importance then, other things being equal, it is difficult to see how it could be permissible to cause Ann to exist.

Another class of views which can be dismissed are versions of Modal Partiality which was discussed above. According to these views, only harms and benefits to people with modal property M have moral significance. If only harms and benefits to, say, actual people matter then it may be impermissible to create Carl in the inverse non-identity case, but it is unclear why it would be permissible to create Ann in the ordinary non-identity case.

In order to avoid Symmetry we therefore need a different view about the moral significance of harms and benefits. The following sections explore what this difference could be.

5 Harming and wrongdoing

An easy way to achieve a deontic asymmetry between the ordinary and inverse Non-Identity Case is to say that harms have moral sig-

nificance and benefits do 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 original Non-Identity Case, were they to be caused to exist. This difference in harming is, on the view under consideration, what explains the normative difference between the two non-identity cases.

A view along these lines is suggested by Boonin (2014, ch. 7). Boonin argues that we should accept 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 not morally wrong.

Together with a counterfactual analysis of harm,¹⁶ these two principles entail that it is not wrong to cause Ann to exist and that it is not wrong to cause Beth to exist. However, this Boonian view does not entail that it is wrong to cause either Carl or Dave to exist in the inverse non-identity case. The view does therefore not show that there is a difference between the two non-identity cases. In order to yield the desired result in the inverse case it has to be assumed that harming, at least under certain circumstances, is sufficient for wrongdoing. A problem for this version of Boonin's view is that it seems to imply that the Inverse Non-Identity Case is a moral dilemma since the couple will end up harming someone, and therefore doing something wrong, whatever they do. However, it does not seem plausible that the inverse cases are moral dilemmas.

To solve this problem we can formulate a version of Boonin's view in the following way:

¹⁶ Boonin does not explicitly endorse the counterfactual analysis as I formulate it here. He merely assumes that making people worse off than they would otherwise have been is *necessary* for harming, not that it is sufficient. However, the stronger version of the counterfactual analysis is, as I will argue, required in order to defend a rejection of Symmetry.

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 wrong iff 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.

By appealing to degrees of wronging, where the degree to which a person is wronged is a function of the degree to which she is harmed, this version of Boonin's view seems able to explain why it is wrong to cause Carl to exist but why it is at the same time permissible to cause both Ann and Beth to exist.

This Boonian view 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, the Boonian view seems superior to the simple axiological asymmetry since it is consistent with Pareto-like considerations. To illustrate, the Boonian view implies that if one alternative is worse for someone, and as least as good for everyone else, then it is wrong. This means that in *same-people cases*—cases where the same people will exist whatever one does—the Boonian view implies that there is no morally relevant difference between cases where we cause a person with a life worth not living to exist and causing a person with a life well worth living to exist. For example, in same-people versions of the two non-identity cases, the Boonian view implies that it is wrong to cause the person with less well-being to exist. It is only in *different people* cases that the Boonian view implies that there is a difference between causing people with lives worth not living to exist and causing people with lives well worth living to exist.

The Boonian view could be criticised on the grounds that it is implausible to hold that only harms, and not benefits, matter morally.

However, since the view implies that it is wrong to cause the person with less well-being to exist in some people cases it might be difficult to find counterexamples which are not non-identity cases. To illustrate, the Boonian view implies that it is not wrong to cause one person with low, but positive, well-being to exist rather than to cause a billion people with very high well-being to exist. Though this seems counter-intuitive a defender of the Boonian view can point out that this is just a version of the ordinary non-identity case. Since the defender of the Boonian view is already committed to the view that it is not wrong in ordinary non-identity cases to cause the person with less well-being to exist, it would be somewhat question-begging to complain that the view implies that it is not wrong to cause the person with less well-being to exist in this case.

Another problem for the view, which does not run the risk of begging the question, can be identified if we consider what the Boonian view identifies as the wrong-maker of certain acts. On the Boonian view, if an act is wrong it is because there is an alternative act such that performing that act would be better for someone. This seems plausible in many cases. For example, an act which causes suffering or which fails to prevent suffering is plausibly wrong because not performing these act would have been better for someone than if the act had not been performed. These are cases where it is plausible to hold that an act is wrong because it harms a person. However, the Boonian view also implies that in cases where it is plausible to hold that an act is wrong because it fails to benefit a person, that act is wrong because it harms that person. This is implausible. For example, if I fail to bestow some good unto some person, and this is wrong, then it is strikingly odd to say that it is wrong because if I were to not bestow this good to this person then I would harm him or her.¹⁷ A more plausible wrong-maker for my action is that I

¹⁷ For example, suppose that I can go to visit my family over the weekend or stay at home. Visiting my family would make them slightly better-off (because we haven't seen each other in a while) but they would not be badly off in any way if I were to stay home. The Boonian view implies, implausibly, that I would harm my family if I decide to stay home.

could benefit the person, but failed to do so.

It might be argued that the Boonian view is compatible with the view that an act can be wrong in virtue of it failing to benefit someone in the following way. It could be claimed that harm and benefit are merely “two sides of the same coin”, and that the claim that an act is wrong in virtue of it harming a person is equivalent to it being wrong in virtue of it failing to benefit that person.¹⁸ The reason for thinking this undoubtedly derives from the counterfactual analysis of harm. Assuming this analysis it is plausible to accept an analogous analysis of benefit: an event benefits a person only if there is an alternative to that act which would be worse for that person. For example, if Ann beats up Beth and Beth would have been better off had Ann not beaten her up, then Ann’s beating of Beth harms Beth. However, because Beth would have been better off had Ann not beaten her up it follows, assuming a counterfactual analysis of benefit, that Ann not beating up Beth benefits her. In this case it might then be plausible to say that there is only one fundamental wrong-maker; namely, that Beth is worse off than she would otherwise have been had Ann not beaten her up.

However, this argument is inconsistent with a plausible extension of the existence requirement. Above I argued that the Boonian view should include as a necessary condition for harming that an event harms a person only if the event occurs and the person exist. The motivation for this requirement was that otherwise the counterfactual analysis of harm implies that a person can be harmed even though she does not exist. It was also noted that without the existence requirement the Boonian view implies that a person who would have a life worth living, were she to exist, is harmed by being caused not to exist. This is a consequence which the Boonian view has to avoid if it is to avoid Symmetry.

Given these remarks about harm it seems plausible to extend the existence requirement to benefits as well. That is, an event benefits

¹⁸ A view along these lines is suggested by Roberts (2011).

a person only if the event occurs and the person exists. However, the “two sides of the same coin” argument implies that this is false. Consider the Inverse Non-Identity Case. In this case, the Boonian view 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, then we would have to say that causing Carl not to exist benefits him. However, this violates the extended existence requirement.

To further illustrate the point that the Boonian view cannot claim that harms and benefits are merely two sides of the same coin, it is useful to consider what the view identifies as the wrongmaker of acts that are wrong in virtue of doing harm. On the Boonian view, if an act ϕ is wrong then the wrong-maker of ϕ is:

Wrong-making-1: ϕ would be worse for p than ψ would be and p would exist if ϕ were performed.

This is a plausible wrong-maker for acts which are wrong in virtue of them harming someone, given a counterfactual analysis of harm. Now, if harms and benefits are merely two sides of the same coin then Wrong-making-1 should also be the wrong-maker for acts which are wrong in virtue of failing to benefit a person. However, this is not the case. Recall that, on the counterfactual analysis, an act harms a person iff that person would have been better off had the act not been performed, *and* the person would exist were the act to be performed. The latter part of this view—the existence requirement—was necessary to add in order to avoid the conclusion that non-existent people can be harmed. However, such a condition is just as plausible to add to an analysis of benefit. That is, an act benefits a person iff it would have been worse for the person, had the act not been performed, and the person would exist were the act to be performed. Therefore, the plausible wrong-maker in cases of failing to benefit is the following:

Wrong-making-2: ϕ would be worse for p than ψ would be and p would exist if ψ were performed.

Wrong-making-1 and Wrong-making-2 are clearly not the same. Therefore, that in virtue of which harming a person is wrong is not equivalent to that in virtue of which failing to benefit is wrong. Note also that a defender of the Boonian view cannot accept that acts are wrong in virtue of Wrong-making-2 and still reject Symmetry. The only alternative for a defender of the Boonian view is to reject the claim that acts can be wrong in virtue of Wrong-making-2. This means that the view is committed to false claims about the moral relevance of benefits. In particular, the “two sides of the same coin” reply is committed to an analysis of benefits which implies that an act can be permissible because it benefits non-existent people. This way of explaining the difference between the cases should therefore be rejected.

6 Conditional beneficence

Above I argued 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 implausible claims concerning benefits. A more plausible view would be one according to which harms always matter morally but the moral importance of benefits is conditional on avoiding harm. Shiffrin (1999, p. 124), for example, makes a distinction between “pure benefits” which she defines as “those benefits that are just goods and which are not also removals from or preventions of harm” and benefits which are also removals or preventions from harm. Shiffrin then argues that pure benefits lack the normative significance which benefiting someone by removing or preventing a harm has. On this “conditional view”, harming a person always has moral significance but benefiting a person has moral significance only if failing to benefit this person would be to do harm.

This view requires some clarification. One complicating factor concerns the distinction between all-things-considered and partial

harm (benefit). The counterfactual analysis of harm described above is plausible read as an analysis of harm all things considered. On this analysis, an event harms a person all things considered iff the person would have been better off, all things considered, had the event not occurred. Partial harm, on the other hand, is often glossed as making the person worse off in some way or respect. To illustrate, a person who undergoes chemotherapy for some particular form of cancer can be harmed partially (because it is so painful) but benefited all things considered (because the benefits of getting rid of the cancer are greater than the painfulness of the treatment).

What the conditional view implies may well depend on whether it is understood in terms of all things considered or partial harm. However, to test whether the view implies that there is a difference between ordinary and inverse non-identity cases we can consider versions of these cases where the all things considered and partial view coincide. That is, we should imagine, somewhat unrealistically, that Ann and Beth in the ordinary non-identity case are not partially harmed in any way. Likewise, we should imagine that Carl and Dave in the inverse non-identity case are not benefited in any way.

A further clarification concerns whether the conditional view should be understood as having an implicit person-affecting restriction. One way to understand the conditional view is that it is necessary for a benefit to have moral significance that *the same person* would have been harmed if she had not been benefited. This version of the view can be formulated more precisely as follows:

Person-Affecting Conditional View: If an act ϕ harms a person p then this is a (moral) reason against ϕ -ing. If an act ϕ benefits a person p then this is a (moral) reason to ϕ only if not ϕ -ing would harm p .

Stating the conditional view in this way captures the idea that harming a person has a certain moral significance by being (moral) reasons. Benefiting a person, on the other hand, are only (moral) rea-

sons on the condition that the person would have been harmed had she not been benefited. The use of 'reasons' here is merely one way of capturing the difference in moral significance between harming and benefiting and is not essential to the view.

This version of the conditional view has several advantages. First, it is consistent with Pareto-like considerations for some people cases. That is, as long as the same people will exist whatever we do, the balance of (moral) reasons will always favour an act which benefits someone and which does not harm anyone else. Second, the view can explain the difference between the two non-identity cases, with some additional assumptions, in the following way. Causing the worse-off person (Carl) to exist in the Inverse Non-Identity Case would harm that person more than causing the better off person (Dave) to exist would harm that person, and both these harms are reasons against causing either person to exist. However, since Carl would be harmed more than Dave, it seems plausible that the reason against causing Carl to exist is weaker than the reason against causing Dave to exist. It therefore seems plausible to conclude that it is not permissible to cause Carl to exist and permissible to cause Dave to exist.

In the original Non-Identity Case on the other hand it is not the case that Ann, or Beth, would have been harmed had they not been caused to exist. That Ann would benefit by being caused to exist is therefore not a reason to cause Ann to exist, and the same goes for Beth. Given that there are no other morally relevant considerations in this case it therefore seems plausible to conclude that it is permissible to create Ann or Beth.¹⁹

The implicit person-affecting component of this view leads to

¹⁹ A possible objection to this view, raised by Bradley (2013), is that on the conditional view the moral significance of benefits depend on what the alternatives are. Suppose for example that we add a third alternative to the ordinary non-identity case where both Ann and Beth exist with neutral well-being (their lives are equally as good for them as non-existence). In this version of the case causing Ann to exist with a life worth living *is* a reason to cause her to exist, since there is an alternative which would harm her (the one where she has a neutral life). I will not pursue this objection further since there are other, more straightforward problems with the conditional view which I discuss below.

some problematic consequences. Consider the following case:

$$\begin{aligned}\phi &= (1, -1, -, -) \\ \neg\phi &= (-, -, 10, -1)\end{aligned}$$

In this case we can either cause two people to exist, one who would have a life barely worth living (1) and one who would have a life slightly worth not living (-1). Alternatively, we can cause two other people to exist, one who would have a life well worth living (10) and one who would have a life slightly worth not living (-1). The Person-Affecting Conditional View is indifferent between these two alternatives, but it seems odd to hold that it is permissible to do ϕ in this case. The oddness is, at least in part, that this case illustrates that the Person-Affecting Conditional View does not fully capture the intuitive idea that benefiting matters conditionally. A part of this intuitive idea is, I think, that benefits are at least tie-breakers in the balance of reasons. That is, if you are going to harm someone (to an equal extent) whatever you do, then you should do that which will benefit people most.

It might be objected that this example begs the question against the Person-Affecting Conditional View since a defender of this view is obviously willing to accept that it is permissible to cause either Ann or Beth to exist in the Non-Identity Case. According to this objection, saying that it is not permissible to ϕ in the case above is tantamount to saying that it is not permissible to cause Ann to exist in the Non-Identity Case, and can therefore not be assumed in this dialectical context. To reply, it is dubious at best, and most probably false, that the two cases are sufficiently similar that it is question-begging to assert that it is not permissible to ϕ in the case just described. The Non-Identity Case involves no harms, while this case does, and what this case brings out is that the Person-Affecting Conditional View assigns too little moral importance to benefits.

One way to refine the conditional view in order to avoid this problem is to drop the person-affecting component:

Impartial Conditional View: If an act ϕ harms a person p then this is a reason against ϕ -ing. If an act ϕ benefits a person p then this is a reason to ϕ only if not ϕ -ing would harm someone.

On this view, we should not be indifferent in the case described above. If we were to ϕ , then we would benefit one person and this has moral significance since if we were not to ϕ then we would harm someone, albeit a different person than the one who would be benefited were we to ϕ . This view also implies, with some further assumptions, that it is permissible to create either Ann or Beth in the Non-Identity Case but that it is impermissible to create Carl in the Inverse Non-Identity Case. The Impartial Conditional View therefore captures the idea that an agent is permissible to do whatever she wants as long as none of her alternatives do harm.²⁰

However, the impartial view has at least as counter-intuitive consequences as the person-affecting view. Consider the following case:

$$\begin{aligned}\phi &= (1, 0, -, -) \\ \neg\phi &= (-, -, 10, -1)\end{aligned}$$

In this case the impartial view implies that the positive well-being which one person would enjoy if we were to ϕ has moral significance since someone would be harmed if we were to not ϕ . The positive well-being which one person would enjoy if we were to not ϕ , on the other hand, does not have moral significance since no-one would be harmed if we were to ϕ . The Impartial Conditional View therefore strongly suggests that it is impermissible to not ϕ .

Suppose however that we grant that it is defensible to hold that not ϕ -ing is permissible, perhaps because one attaches a much greater weight to harming than to benefiting when balancing reasons. However, the example just described still has force against the Impartial, and the Person-Affecting, Conditional View. What is so strikingly

²⁰ See for example Scheffler (2003, pp. 182-4) who suggests that agents have a “no-harm prerogative”.

odd about these views is that they imply that not ϕ -ing does not have anything going for it; there is *no* reason to not ϕ . There simply is no reason which favours not ϕ -ing which can be balanced against the reasons against not ϕ -ing. A more plausible approach to both cases is that benefiting has *some* moral significance, though perhaps not equal to the significance of harming. This approach can easily explain why it is permissible, and perhaps required, to not ϕ in both cases. However, it is obvious that this view implies Symmetry.

A defender of the Impartial Conditional View might admit that it has counter-intuitive consequences but insist that accepting Symmetry is even more counter-intuitive. The Impartial Conditional View seems to be a coherent way of explaining the difference between the two non-identity cases and if one's commitment to there being such a difference is stronger than one's belief that it is permissible to not ϕ in the cases described above, then this may be offered as an argument for accepting these counter-intuitive conclusions.²¹

It is not easy to assess the claim that the Impartial Conditional View is more plausible—or rather, less implausible—than Symmetry. A noteworthy difference between the two is that the Symmetry is a very restricted claim which does not have very widespread consequences. The Impartial Conditional View, on the other hand, has very far-reaching consequences. For example, the Impartial Conditional View clearly has consequences for what we should do in *different number cases*; cases where our acts affect not only the well-being and identity of those who will exist as a consequence, but also the number of people who will exist. In particular, the Impartial (and the Person-Affecting) Conditional View implies the repugnant conclusion.²² Symmetry, on the other hand, has no implications in different-number cases. These considerations suggest that it is far from clear that the Conditional View has an advantage

²¹ This strategy is employed by Benatar (2006, p. 32) and, to some extent, by Boonin (2014, ch. 7).

²² The repugnant conclusion is usually stated in axiological terms. For a discussion of deontic versions of the repugnant conclusion, see Arrhenius (2000, ch. 11).

over Symmetry when it comes to intuitiveness.

7 Concluding remarks

Some philosophers have argued that it is sometimes permissible to cause either person to exist in ordinary non-identity cases. However, no one, I believe, holds that it is permissible to cause either person to exist in inverse non-identity cases. If one holds these two beliefs then one rejects Symmetry: it is permissible for the couple to have a child now in the ordinary Non-Identity Case iff it is permissible for the couple to have a child now in the Inverse Non-Identity Case. In this paper I 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 versions of the non-identity problem. In particular, I have shown that an axiology which is “modally partial”—it discounts the well-being of people with a certain modal property—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. I outlined two such views on which Symmetry is false: the person-affecting conditional view and the impartial conditional view. The characteristic feature of both these views is that benefits matter, morally, only on the condition that they avoid harms and this feature allows both views to reject Symmetry. A view which merely assigns different moral weights to harms and benefits will, I argued, imply Symmetry. However, making the moral relevance of benefits conditional on avoiding harm has counter-intuitive consequences which should make us doubt that these views are plausible alternatives to Symmetry.

References

- Arrhenius, G. (2000), *Future Generations: A challenge for moral theory (Dissertation)*, Uppsala: Acta Universitatis Upsaliensis.
- Arrhenius, G. & Rabinowicz, W. (2010), Better to be than not to be, in H. Joas & B. Klein, eds, 'The Benefit of Broad Horizons', Leiden: Brill.
- Belshaw, C. (2003), 'More lives, better lives', *Ethical Theory and Moral Practice* **6**, 127–141.
- Benatar, D. (2006), *Better Never to Have Been: the Harm of Coming Into Existence*, Oxford: Clarendon Press.
- Bigelow, J. & Pargetter, R. (1988), 'Morality, potential persons and abortion', *American Philosophical Quarterly* **25**(2), 173–181.
- Boonin, D. (2014), *The non-identity problem and the ethics of future people*, Oxford: Oxford University Press.
- Bradley, B. (2013), 'Asymmetries in benefiting, harming and creating', *Journal of Ethics* **17**(1-2), 37–49.
- Broome, J. (2004), *Weighing Lives*, Oxford: Oxford University Press.
- Brown, C. (2011), 'Better never to have been believed: Benatar on the harm of existence', *Economics and Philosophy* **27**, 45–52.
- Feinberg, J. (1986), 'Wrongful life and the counterfactual element in harming', *Social Philosophy and Policy* **4**(1), 145–178.
- Holtug, N. (2010), *Persons, Interests, and Justice*, Oxford: Oxford University Press.
- Jackson, F. & Pargetter, R. (1986), 'Oughts, options, and actualism', *The Philosophical Review* **95**(2), 233–255.
- Johansson, J. (2010), 'Being and betterness', *Utilitas* **22**(3), 285–302.
- McMahan, J. (2009), Asymmetries in the morality of causing people to exist, in M. A. Roberts & D. T. Wasserman, eds, 'Harming Future Persons', Dordrecht: Springer.
- Österberg, J. (1996), Value and existence: The problem of future generations, in S. Lindström, R. Sliwinski & J. Österberg, eds, 'Odds and Ends', Uppsala: Department of Philosophy, Uppsala University.
- Parfit, D. (1984), *Reasons and Persons*, Oxford: Oxford University Press.
- Parsons, J. (2002), 'Axiological actualism', *Australasian Journal of Philosophy* **80**(2), 137–147.

- Rabinowicz, W. (2009), 'Broome and the intuition of neutrality', *Philosophical Issues* **19**(1), 389–411.
- Roberts, M. A. (2003), 'Is the person-affecting intuition paradoxical?', *Theory and Decision* **55**, 1–44.
- Roberts, M. A. (2007), 'The non-identity fallacy: harm, probability and another look at Parfit's depletion example', *Utilitas* **19**(3), 267–311.
- Roberts, M. A. (2011), 'The asymmetry: a solution', *Theoria* **77**, 333–367.
- Savulescu, J. (2001), 'Procreative beneficence: why we should select the best children', *Bioethics* **15**(5/6), 413–426.
- Scheffler, S. (2003), *The Rejection of Consequentialism (Revised edition)*, Oxford: Clarendon Press.
- Shiffrin, S. (1999), 'Wrongful life, procreative responsibility, and the significance of harm', *Legal Theory* **5**, 117–148.
- Singer, P. (1993), *Practical Ethics*, Cambridge: Cambridge University Press.