

**Manuscript version: Author's Accepted Manuscript**

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

**Persistent WRAP URL:**

<http://wrap.warwick.ac.uk/165550>

**How to cite:**

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: [wrap@warwick.ac.uk](mailto:wrap@warwick.ac.uk).

# The Impure Non-Identity Problem\*

*Patrick Tomlin*

Final version published in in Jeff McMahan, Tim Campbell, James Goodrich, and Ketan Ramakrishnan eds., *Ethics and Existence: The Legacy of Derek Parfit* (Oxford: Oxford University Press, 2022)

## I. Introduction

The aims of this chapter are to convince you of the following claims and their importance, and to explore their implications for moral theory. First:

The Impure Non-Identity Claim: a set of acts can affect who comes into existence and be neither harmful to nor worse for any particular individual, even though each act in the set does harm, in the sense of being worse for, some individual or individuals.

Given this, I will argue for this second claim:

Two Non-Identity Problems: Impure Non-Identity cases pose important moral questions which Pure Non-Identity cases do not. Therefore, there are two importantly different non-identity problems.

---

\* I am grateful for discussion with the attendees of the conference held in Derek Parfit's memory at Oxford in April 2018, especially Mike Otsuka and Jonathan Dancy. I am grateful to Jeff McMahan and Victor Tadros for discussion, extensive written feedback, and conflicting advice.

Before presenting my arguments for these claims, and discussing the moral issues they raise, I will, by way of background, describe two issues with which the chapter engages. The philosophical problems I investigate in this chapter sit at the intersection of these two issues, both of which Derek Parfit illuminated in more-than-seminal ways.

The first is the non-identity problem, which I will discuss in more depth in the main body of the chapter. One of my arguments here is that it is really two problems, or that there are two sub-problems: there are Pure Non-Identity Cases, which raise a set of issues, and Impure Non-Identity Cases, which raise additional issues.

Here is a Pure Non-Identity Case: A couple planning to have children find out that, due to a virus one of them currently has, if they conceive within the next month, the baby will have some serious health issue that will negatively affect its quality of life. Its life will nonetheless be worth living. If the couple wait, however, their baby will not have this serious health issue. On the one hand, it seems the couple ought to wait because of the way the health issue will affect the child's wellbeing. On the other, if they wait, they will have a different child. Having the child earlier, with the attending health problem, is not *worse for* that child – the only other option is that the child would never exist.

The non-identity problem is the problem whether we can defend and explain the intuition that the couple who don't wait act wrongly. If they do act wrongly, this seems to be a major challenge for person-affecting or complaint-based theories of

morality. One of Parfit's many towering achievements was to introduce us to this problem, and, perhaps more importantly, to show us how far it reaches in its challenges to all moral theories, and what its implications are for both public policy and private decisions. It is a foundational element of perhaps the most revolutionary part of the most revolutionary book in twentieth century moral philosophy.<sup>1</sup>

The second issue with which this chapter engages is what Parfit called 'moral mathematics': puzzles that arise when we collect acts together into courses of action, either as groups or as individuals. Parfit made great strides in exposing these puzzles, and investigating them, in his 'Five Mistakes in Moral Mathematics'.<sup>2</sup> One such puzzle involves overdetermined harm – groups of acts in which we appear to collectively harm somebody, but in which no individual act appears to harm. Consider a course of action in which five individuals simultaneously stab a person, for example the murder of Julius Caesar. Collectively, they kill Caesar. But had each of them not acted as he did, Caesar would still have died – one stab is sufficient to kill him. Had one individual, Brutus, not stabbed him, Caesar would still have been killed. And so Brutus does not appear to kill Caesar. The same applies to all the other individuals. Collectively they kill Caesar, but no individual kills him. If harming is making people worse off than if you had not acted, no individual harms Caesar.

---

<sup>1</sup> Parfit, *Reasons and Persons*, Part IV.

<sup>2</sup> *Ibid.*, ch. 3.

The cases I want to consider here are, in some senses, the opposite of cases of overdetermined harm: in my cases, each individual act makes someone worse off, but the collection of acts does not make anyone worse off.

## II. Two Non-Identity Cases

The non-identity problem arises when an agent faces a choice of whether to bring into existence a person who will be well off, or to bring into existence a different person who will be less well off. On the one hand, it seems intuitive there's at least some moral reason to bring into existence the well off person. On the other, since one is choosing which of two people (or groups of people) to bring into existence, neither choice is *worse for* or *better for* any particular person (assuming all will have lives worth living). Since nobody is made worse off, arguably, nobody has any complaint about the choice.

For the sake of clarity and precision, it will help to introduce some terminology:

*Harm*: someone is made worse off than they otherwise would have been.<sup>3</sup>

*Impersonal Loss*: someone is caused to exist who will have lower wellbeing than some other person who otherwise would have existed.<sup>4</sup>

---

<sup>3</sup> There are, of course, all sorts of objections to counterfactual views of harm, especially simple counterfactual views. These stem from overdetermination and pre-emption cases in particular. But this definition of harm nevertheless seems to capture what is at stake in non-identity cases. In cases in which identity is not affected, someone has a complaint because they have been made worse off. In non-identity cases, someone is badly off but they have not been made worse off.

<sup>4</sup> In Parfit's later work on the non-identity problem, he affirmed a wide person-affecting view to underpin his No Difference View. Under the wide person-affecting view, impersonal loss would better be defined as a lower level of existential benefit. I do not think that this difference is important for the issues and arguments presented here. For the wide person-affecting view, see: Parfit, 'Future People, the Non-Identity Problem, and Person-Affecting Principles.'

*Negative Wellbeing Effect*: an umbrella term encapsulating both harm and impersonal loss – where wellbeing is lower than it otherwise would have been, regardless of whether the same person, or some different person, would have otherwise existed (keeping numbers of people constant).

*Identity-Fixing Act*: The act of selecting an option which determines who will exist.

Non-identity cases are, therefore, those in which an option which will have a negative wellbeing effect will also be an identity-fixing act. Since both of these features are present, this means that the negative wellbeing effect is one of impersonal loss rather than harm.

Here are two famous cases that Parfit uses to illustrate the non-identity problem. They demonstrate its reach by showing how it applies to both personal and political decisions:

*The 14-Year-Old Girl*: This girl chooses to have a child. Because she is so young, she gives her child a bad start in life. Though this will have bad effects throughout this child's life, his life will, predictably, be worth living. If this girl had waited for several years, she would have had a different child, to whom she would have given a better start in life.<sup>5</sup>

---

<sup>5</sup> Parfit, *Reasons and Persons*, p. 358.

*Depletion*: As a community, we must choose whether to deplete or conserve certain kinds of resources. If we choose Depletion, the quality of life over the next three centuries would be slightly higher than it would have been if we had chosen Conservation. But it would later, for many centuries, be much lower than it would have been if we had chosen Conservation. It is not true that, whichever policy we choose, the same particular people will exist in the further future. Since the choice between our two policies would affect the timing of later conceptions, some of the people who are later born would owe their existence to our choice of one of the two policies.<sup>6</sup>

Parfit appears to see these two cases as posing the same problem, and the vast literature that has followed in his wake has tended to agree.<sup>7</sup> But they're importantly different, as Depletion potentially poses moral problems that 14-Year-Old Girl does not.

This is because, in my terminology, 14-Year-Old Girl is a Pure Non-Identity Case while Depletion is potentially an Impure Non-Identity Case. Pure Non-Identity Cases are those in which each identity fixing act also causes impersonal loss. There is, in other words, no way that the resulting person or persons could have existed at a higher level of wellbeing. Impure Non-Identity Cases, in contrast, are those in which a course or

---

<sup>6</sup> Ibid., pp. 361-362. All the words in the case are Parfit's from across these two pages, though I have reordered some passages.

<sup>7</sup> See, for example, the way that they are initially introduced and discussed in Roberts., 'The Nonidentity Problem.' However, both John Broome and Michael Otsuka describe climate change policy as an area where the policy decisions are non-identity cases but individual emissions decisions are not. Otsuka, 'How it makes a moral difference'; Broome, *Climate Matters*, p. 64.

collection of acts is both identity-fixing and causes impersonal loss, but different acts within the course or collection have the identity-fixing and negative wellbeing impacts for given individuals. This means that, without some of the acts, the individual could have existed at a higher level of wellbeing. As a result, the collection of acts cause impersonal loss, but the individual acts are harmful.

In order for a case to be an Impure Non-Identity Case, it must be a course or collection of acts. But this is not sufficient. Cases involving a plurality of acts can still be Pure Non-Identity Cases. Consider this case:

*One Thousand Acts:* Government is considering funding an education policy that will convince teenagers to delay parenthood. If they fund the policy, one thousand teenagers who would otherwise have conceived will choose to delay parenthood.<sup>8</sup>

This is a policy-level decision, which involves affecting the choices and actions of many others. But it is still a straightforward non-identity case – it is simply *14-Year-Old Girl* writ large. Either the teenagers have children, and one thousand children will exist, with less good lives, or they wait and some other one thousand children will exist, and will have better lives than the teenagers' children would have had. If the government don't fund the policy, the children of the teenagers will have no

---

<sup>8</sup> This case is structurally similar to one of the policies that Parfit introduces in his *Two Medical Programmes* case, in order to argue for his No Difference View. Parfit, *Reasons and Persons*, p. 367; Parfit, *On What Matters: Volume Two*, p. 221.



complaint, for if the government had funded the policy they simply wouldn't have existed. The negative wellbeing effect is one of impersonal loss.

In sum, Pure Non-Identity cases are those in which the resulting person or persons can have no complaint about, and have not been harmed by, *either* any individual act, *nor* the policy or collection of acts.

In contrast, Impure Non-Identity cases are those in which, because it is identity-fixing, the collection of acts harms no one, but at least some of the individual acts do harm someone, and perhaps *all of them* do. I will focus on this latter kind of case, in which each and every individual act harms someone, even though the collection of acts harms no one.

### III. Impure Non-Identity: Circle Cases

It will be easiest to understand how Impure Non-Identity Cases work by looking at some of them. There are at least two kinds of Impure Non-Identity case. In this section, we will examine the general structure of one of these. In order to do so, I need to introduce a further bit of terminology. Imagine that my parents could either have had me, or waited a month and had some other child. That other child – let's call her Patricia – and I are, obviously, different individuals. But let's say my parents would only ever have had one of us. Therefore, Patricia and I share a *placeholder*. There is some space in the world for my parents' child – that is the placeholder – and it could have been filled by either Patricia or me. In the following, let Child A etc. stand for the

placeholder, and Child A<sub>1</sub>, Child A<sub>2</sub> etc. stand for the different individuals who might have filled that placeholder.

*Circle*

Act 1 causes Child A<sub>1</sub> to exist instead of Child A<sub>2</sub>, and has a negative wellbeing effect on Child B.

Act 2 causes Child B<sub>1</sub> to exist instead of Child B<sub>2</sub>, and has a negative wellbeing effect on Child C.

Act 3 causes Child C<sub>1</sub> to exist instead of Child C<sub>2</sub>, and has a negative wellbeing effect on Child D...

Act N causes Child X<sub>1</sub> to exist instead of Child X<sub>2</sub>, and has a negative wellbeing effect on Child A.

Each act harms someone. For example, consider Act 2. If Act 2 does not occur, then there will be no negative wellbeing effect on the Child C placeholder. Let's say that Act 3 goes ahead regardless of whether Act 2 does. Then C<sub>1</sub> will exist, and there will be no negative wellbeing effect on C<sub>1</sub>. So, Act 2 makes a difference to whether Child C has a better or worse life. But since C<sub>1</sub> exists regardless of whether Act 2 goes ahead or not (it is Act 3 that is identity-fixing for Child C), Act 2 *harms* C<sub>1</sub>.

However, this *course* of action does not harm anybody. C<sub>1</sub> is harmed by Act 2. But without the course of action, C<sub>1</sub> would not have existed. Therefore, provided C<sub>1</sub> has a life worth living, they are not harmed by the course of action.

Let's call these *Circle Cases*.<sup>9</sup> There are some sub-types of Circle Cases. They differ from one another in two respects. The first concerns whether a single person or a group of people carry out the individual acts. The second concerns the relationship between the group of acts and how coordinated they are. One especially interesting kind of case is where a group or individual (a policy maker, for example) makes a decision (or an order) that they foresee will cause others to act in certain ways.

Here's a case in which a group of agents carries out the acts, prompted by the decision of a policy maker:

*Pollution Regulation:* A government official must decide whether to relax pollution restrictions on factories. At present, factories are required to employ pollution monitors. Companies will be free to fire these monitors if pollution restrictions are relaxed. If the restrictions are relaxed, the official foresees the following chain of acts will occur:

Act 1: Company A will fire Monitor A. Monitor A will get a better job, and she and her partner will move to live by Lake B. They will have child A<sub>1</sub>, instead of Child A<sub>2</sub>. Increased pollution in Lake A caused by Company A will negatively affect any child who grows up next to Lake A.

Act 2: Company B will fire Monitor B. Monitor B will get a better job, and he and his partner will move to live by Lake A. They will have child B<sub>1</sub>, instead of Child B<sub>2</sub>. Increased pollution in Lake B caused by Company B will negatively affect any child who grows up next to Lake B.

---

<sup>9</sup> I call them this because Act N ends up having a negative wellbeing effect on Child A, and so the chain loops back round. We can also imagine chain cases without this looping.

First let's examine the acts of the two companies. In firing their pollution monitor, and increasing pollution in Lake B, Company B have a negative wellbeing effect on any child who grows up in next to Lake B. Due to Act 1, this is Child A<sub>1</sub>. However, since Company B's decision has no effect on whether Child A is Child A<sub>1</sub> or A<sub>2</sub>, they harm Child A<sub>1</sub> through their decision. Company A's act has a negative wellbeing effect on any child who grows up next to Lake A. Due to Act 2, this is Child B<sub>1</sub>. But since Company A's acts have no impact on Child B's identity, this is a harm to Child B<sub>1</sub>. Act 1 and Act 2 are therefore harmful acts.

Now let's examine the policy-maker's act. The policy-maker's act – namely instituting the policy – will harm nobody. If they do not act, A<sub>2</sub> and B<sub>2</sub> will exist. If they do act, there will be a negative wellbeing effect, but it will be an impersonal loss, as A<sub>1</sub> and B<sub>1</sub> will exist, and will have lives worth living.

In making her decision about whether or not to relax pollution controls, the policy-maker thus finds both of these statements to be true:

*Statement 1:* If I relax pollution controls, two children will be harmed by polluters.

*Statement 2:* If I relax pollution controls, I will cause impersonal loss, but I will not cause any harm.

In *Pollution Regulation* a policy-maker makes a decision, and then independent agents implement, or react to, that decision with their own decisions. However, imagine a

modified version of the case in which a single company owns two factories, A and B. The head of the company must decide whether to pursue a policy of firing pollution monitors or not. If she does fire the monitors, she foresees the effects will be the same as those in *Pollution Regulation*. In considering whether or not to fire the monitors, she finds both the following statements to be true:

*Statement 1:* If I decide to relax pollution controls, I will harm two children.

*Statement 2:* If I decide to relax pollution controls, I will harm nobody.

Statement 1 is true of individual acts she will undertake if she decides to implement the policy. Statement 2 is true of the collection of acts she will undertake if she decides to implement the policy.

In Pure Non-Identity Cases, individual acts have a negative wellbeing effect on, and are identity-fixing for, the *same* individuals. Therefore, the individual acts each produce impersonal loss, but no harm. In these Circle Cases, all the acts are both identity-fixing *and* produce negative well-being effects, but different acts within the group of acts have the negative wellbeing effect and are identity-fixing for, given individuals. Therefore, there was a chance (where the acts are sufficiently independent) that the individual could have existed but without the negative wellbeing effect. So the negative wellbeing effect of the act is a harm. But the collection of acts, as it has both a negative wellbeing effect *and* is identity-fixing, is a case of impersonal loss.

#### IV. Impure Non-Identity: Overdetermination of Identity

In this section, I want to discuss a different set of Impure Non-Identity cases. These are courses of action that have negative wellbeing effects *and* are identity-fixing, but in which no individual act from the set determines identities, since the identities are overdetermined within the set of acts. Here's an attempt at formulating the general structure of such a case.

*Overdetermination of Identity:* If N or more acts are performed, Group A will exist. If fewer than N acts are performed, Group B will exist. A set of acts, each of which later negatively affects an individual, is performed. There are more than N acts.

Here's a more concrete example.

*Landmines:* The commanding officer in a war orders 2000 soldiers lay one landmine each. They all foresee that these will have the intended effect of deterring the enemy. They also foresee, however, that many years later, each of these mines will be stepped on by a child, so that 2000 children will be seriously wounded. If, as in fact happens, 1000 or more landmines are laid, families will vacate the area, and Set A of 2000 children will be born. If fewer than 1000 landmines had been laid, the families would have stayed, and Set B of 2000 children would have been born. Later the 2000 children of Set A are seriously injured.

The commander's act of ordering the laying of the mines is non-identity case. The children who are injured, Set A, would not have existed if the course of action ordered by the commander had not taken place. No one, therefore, is harmed by the

commander's action, even though no children would have been injured had he not ordered the laying of the mines. The loss is impersonal. But the act of each individual soldier – each laying of a landmine – is *not* a non-identity case. Each soldier's act, therefore, *harms* a child.

Consider the contrasting ways in which the commanding officer and an individual soldier might assess their acts. The commanding officer might think, of his act of ordering that the landmines be laid, 'it was ordering the laying of these landmines that caused Set A of children to exist. If I hadn't ordered the laying of the landmines, Set B would have existed instead. And so, even though many children were injured, none of them were harmed by my order.' Let's now take a specific soldier, Soldier 1173, and his act of laying a landmine. Soldier 1173's landmine injured Child A1173. But this child would have existed whether Soldier 1173 had laid his landmine or not. While the *course of action* determined the identity of the children who were later injured by the landmines, no *individual act* of laying them did. If Soldier 1173 hadn't laid his landmine, there would still have been more than 1000 landmines laid, and so Set A would still have existed. So, if Soldier 1173 hadn't laid his landmine Child A1173 would have existed, but not have been injured. Therefore, Soldier 1173 harmed Child A1173. The same reasoning applies to each of the 2000 soldiers and the child injured by the landmine they laid.

Imagine that many years later the commanding officer is reflecting on his war experience, and he thinks about this command in particular. He finds both the following statements to be true:

*Statement 1:* I ordered 2000 soldiers to perform acts which, foreseeably, seriously harmed 2000 children.

*Statement 2:* My order did not harm any children.

## V. Moral Implications

We have seen that Pure Non-Identity cases are non-identity cases all the way down. Either they consist of one act (as in *14-Year-Old Girl*), or a series of acts in which each individual act is a non-identity case of its own (as in *One Thousand Acts*). There is unquestionably no harm in such cases (insofar as we understand harm as being made worse off). By contrast, Impure Non-Identity Cases are harmful when viewed as individual acts, but not when viewed as a collection of acts.

The literature on the Non-Identity Problem, and its moral significance, is vast, and I cannot rehearse, or do justice to, all of it here. But, broadly speaking, there are two groups of responses to the problem. The first, endorsed by Parfit<sup>10</sup>, is the No Difference View:

*No Difference View:* It makes no moral difference whether a negative wellbeing effect is an impersonal loss or a harm.

---

<sup>10</sup> Parfit, *Reasons and Persons*, p. 369; *On What Matters: Volume Two* (Oxford: Oxford University Press, 2011), pp. 219-221; 'Future People, the Non-Identity Problem, and Person-Affecting Principles'. Parfit explored two different rationales for accepting the No Difference View – impersonal principles, and wide person-affecting principles.



The second group of responses to such cases (which is a varied group) is what I will call the Difference View. The Difference View, in essence, is any view which rejects the No Difference View, and claims that harm is morally significant:

*Difference View:* It makes a moral difference whether a negative wellbeing effect is an impersonal loss or a harm. Either impersonal losses cannot make actions wrong, or harms are, all else equal, harder to justify than impersonal losses.

As I have stressed, dividing the terrain into these two camps seriously simplifies a complex literature. But doing so is helpful for our purposes here.

Here are two interesting sets of moral questions raised by Impure Non-Identity Cases. The first concerns the general question of how to respond to the non-identity problem – the question of whether we should accept some version of the Difference View, or some version of the No Difference View. I will argue that consideration of Impure Non-Identity Cases should be considered here, and could push us toward the No Difference View.

The other interesting set of questions raised by Impure Non-Identity Cases concern how those who endorse the Difference View ought to morally assess both the policy-maker's acts and the individual acts in such cases. Since, for those who hold the No Difference View, whether a case is a non-identity case or not makes no moral difference, the Impure cases do not pose particularly difficult questions. But if we hold some variant of the Difference View, I will argue, difficult moral questions arise in such cases.

### *Impure Non-Identity Cases and the No Difference View*

That there are Impure Non-Identity cases may affect how plausible we find the No Difference View. In my view, it might push us toward that view.

If we are attracted to the Difference View, we must place a lot of moral weight on the question of whether harm has occurred or not. In Pure Non-Identity Cases whether harm has occurred is clear cut – the negative wellbeing effect is clearly not harm. Therefore, we see a bright line that is capable of bearing the moral weight that the Difference View places on it. But we now see that in these Impure Non-Identity Cases, matters are more complex, for the parties are harmed by individual acts, but not by the collective of acts, or policy. Therefore, whether a person has been harmed is not clear cut – it depends upon whether we are asking the question of the individual acts or of the policy. And whether this harm matters morally will depend upon which we ought to ask the question of. For example, if we should only ask the question of the policy, it might be true that the individual acts harm, but morally this does not matter.

Take *Landmines* as an example. If Child A1173's moral standing to make a complaint relies on whether or not she would have existed *without* the landmines that injured her being laid, then whether or not she has a complaint appears to depend upon whether she is addressing the complaint to the commanding officer or Solider 1173. This may well make us sceptical of putting too much weight on the distinction between harm and impersonal loss – for the same injury can be *both*, depending on whether we are looking at individual acts or collections of acts.

A second way in which Impure Non-Identity cases might push us toward the No Difference view is that whether or not something is a Pure Non-Identity Case can hang on what seem to be morally inconsequential details. For example, compare Landmines with a case in which the commanding officer simply presses a button and 2000 landmines are released (call this case Automated Landmines). This is a single act, and a Pure Non-Identity Case – the same act makes it the case that Set A of children will exist, and that they are later injured, and so nobody is harmed. We might question whether the commander pressing a button versus ordering the soldiers can make a big moral difference, when the same people are created and injured either way. Indeed, imagine the Commanding Officer faces a choice between ordering the soldiers or pressing the button. Would he have moral reason to press the button, since there will then be no harm? What if pressing the button injures an additional child, compared with ordering the soldiers? Would the moral gravity of one child being harmed plus the impersonal loss of two thousand children having less good lives still be the lesser evil compared to two thousand children being harmed, even though the two thousand are the very same children in both cases? That would seem a hard sell to the additional child who is injured.

Some philosophers with whom I have discussed these issues feel that they seriously undermine the Difference View. Others have found their confidence in the Difference View unshaken – after all, they say, we know that collections of acts throw up all kinds of morally knotty problems, as Parfit’s moral mathematics showed us. Why should we expect population ethics to be any different? The challenge for the Difference View is

to address these problems, not to fold in their face. How the Difference View might go about addressing these issues is something I address in the next section.

At the least, however, it does seem that consideration of Impure Non-Identity Cases should be part of the debate. The debate has been constructed around Pure Non-Identity Cases. That there are these more complex Non-Identity Cases should affect the way we address the issues, and test the views. In particular, once we have figured out how the Difference View can best handle such cases, that will in turn affect our assessment of that view.

#### *Impure Non-Identity Cases and the Difference View*

What should proponents of the Difference View say about Impure Non-Identity cases? This is an important question for two reasons. First, many accept some version of the Difference View, and, as I hope to show, the Impure Non-Identity Cases offer difficult questions for the view. Second, how the Difference View handles such cases will help us choose between it and the No Difference View.

We should ask two moral questions of Impure Non-Identity Cases: What should we say, morally speaking, of bringing about the course of conduct (e.g., the policy-maker's, or the Commanding Officer's act)? And what should we say, morally speaking, about the individual acts within the course of conduct?

For simplicity, let's say we're committed to a view in which harms are twice as hard to justify as impersonal losses of the same magnitude. Harms are given a  $2W$  moral weighting, while impersonal losses are given a  $W$  moral weighting.

So far as I can tell, there are, broadly speaking, four options as to how we should treat Impure Non-Identity Cases:

1. The policy-maker's act should be viewed as causing impersonal loss ( $W$ ), while the individual actors' acts should be viewed as causing harm ( $2W$ ).
2. Both policy-maker and individual actors' acts should be viewed as causing impersonal loss ( $W$ ). In other words, the acts of all involved should be treated in the same way as if this were a Pure Non-Identity Case.
3. Both policy-maker and individual actors' acts should be viewed as causing harm ( $2W$ ). In other words, the acts of all involved should be treated in the same way as if this were not a Non-Identity Case at all.
4. All actors' acts should be treated as neither like ordinarily harmful actions or as Pure Non-Identity Cases. Impure Non-Identity Cases, and the moral assessment of them, lie somewhere in between.

The idea in Option 1 is that we recognise that the policy maker does no harm, while the individual actors do. So when we assess the policy maker's act, we assess it as we would a non-identity case, and when we assess the individual actor's acts, we assess them as harmful actions.

This view has two major problems. The first is that, in cases such as the modified version of *Pollution Regulation*, where the same person makes the decision to initiate the course of conduct, and carries out the individual acts, the person is left with no guidance, or conflicting guidance, as to how to assess their behaviour. Qua policy-maker, they are told to treat this as a Non-Identity case, and the ensuing negative wellbeing effects as impersonal losses, whilst qua individual actor, they are told to weight the negative wellbeing effects as harms. But if, according to the Difference View, whether or not the course of conduct is permissible depends on whether the negative wellbeing effects are impersonal losses or harms, then the person is told that the course of conduct is permissible but none of the individual acts are.

The second problem with Option 1 is that it seems to lead to odd situations in cases such as Landmines. In Landmines, according to Option 1, the commanding officer's act of ordering the laying of the landmines ought to be weighted as causing impersonal loss, while the soldiers acts of laying them ought to be weighted as causing harm. Imagine that the country the soldiers fight for is under threat. In order for their acts of war to be just, they must be proportionate. Imagine that, under the Difference View, the goodness of defending the country will render the impersonal loss of 2000 injuries proportionate, but will not render 2000 harms of the same magnitude proportionate.<sup>11</sup> If this is the case, under Option 1, the commander's orders are proportionate, and he is permitted to issue them, but the soldiers' acts are disproportionate, and they are not permitted to carry them out.

---

<sup>11</sup> I explore this case, and its consequences, in greater detail in Tomlin, 'Proportionality in War: Revising Revisionism.'

Options 2 and 3 both insist that the policy-maker's and the individual agents' acts must be assessed in the same way – either as Non-Identity cases, causing impersonal loss, or as causing harm. In other words, according to Option 2, individual agents can either appeal to the wider course of action of which they are part to claim that they are not *harming* those affected by their acts (or at least, morally speaking, that their acts should not be weighted as harmful acts), or the policy-maker must accept that they are causing harm.

It is tempting to think that *either* Option 2 or Option 3 must be the correct one. But this is not necessarily the case.<sup>12</sup> It is possible that sometimes Option 2 is appropriate and sometimes Option 3 is. We should also consider that policy-makers and the acts that are caused by their policies have a variety of differing relationships, and these may affect whether it is appropriate to apply Option 2 or Option 3. In particular, if it is the case that either Option 2 or Option 3 is appropriate, we should consider the nature of the relationship between the policy-maker and the individual actors, and the nature of the relationship between the policy and the individual acts. By way of example, let's compare *Pollution Regulation* and *Landmines*.

In *Pollution Regulation* the relationship between the principal actors is a regulatory one between the government official and companies who trade within that jurisdiction. In *Landmines* the relationship between the actors is one of between

---

<sup>12</sup> I am grateful to Mike Otsuka for useful discussion here.

commanding officer and his troops. Both of these are relationships of authority, though perhaps importantly different ones. Regulators set terms within which companies or individuals choose how to act. Commanding officers direct those under their command.

In terms of the relationship between the policy and the individual acts, the Commanding officer *commands* his troops to lay the landmines, in the expectation that they will do so. That is the desired outcome. In *Pollution Regulation*, meanwhile, the regulator removes a barrier which the companies take advantage of. The government official does not tell them to, or encourage them to, or hope that they will, fire their pollution monitors. He simply foresees that they will do so.

This discussion is possibly driving us toward the vexed issue of group agency – something I cannot tackle here. What I will say is that the more authority the policy-maker has, and the extent to which the acts that cause the negative wellbeing effects are not merely foreseen but intentionally brought about by the policy maker, the more plausible it is that we should see these acts not merely as a collection of independent acts, but as a cohesive group of acts. And then it is more plausible that the policy-maker should see the acts of the individual actors as her ‘own’, and that the individual agents can appeal to the wider context in assessing their acts as causing Impersonal loss rather than harm.

So far we have examined two different factors that may be relevant to whether or not a policy ought to be assessed as causing harm or Impersonal loss. Here is a third: is the



policy designed to produce individual acts, or a collective of acts? Is it that each individual act makes sense, in terms of the aims of the policy-maker, on its own terms, or does it only make sense as part of a collective. Imagine, for example, that a policy-maker creates a policy which aims to incentivise ten actors to each do Act A. Every time act A is performed, it creates 20 units of good. Now imagine a policy-maker creates a policy-maker which aims to incentivise ten actors to each do Act B. If all ten perform the action, 200 units of good will be created. But if fewer than ten actors perform Act B, no good will be created. Both policies are Impersonal Non-Identity Cases.

Morally-speaking, should we treat the negative wellbeing effects as impersonal losses (as per Option 2), or as harms (as per Option 3)? My claim here is that it seems more plausible to treat the case in which Act B is encouraged as impersonal loss. This is because the whole course of action, if it were to be justified, must be justified together – it can only be justified if there is sufficient prospect of all ten performing Act B, for if fewer than ten perform Act B no good will be done. So because the policy's putative justification is as a collection of acts, its costs can also be viewed as a collection of acts. Meanwhile, the policy which seeks to encourage Act A does not need to be justified as a whole. Each Act A is independently producing good (as well as negative wellbeing effects). In seeing whether the goods outweigh the bads, we ought to look at each act independently, which would involve weighing the bads as harms (2W).

I have thus far examined Option 2 and Option 3. I have suggested that we do not need to make a single choice between them, but rather look at several factors – relationship

between policy-maker and actors, relationship between policy and the acts it causes, and the justification of the policy – as to whether Option 2 or Option 3 applies.

However, these considerations may instead push us toward Option 4. Option 4, like Options 2 and 3, suggests that we cannot view the policy-maker's act as one thing and the individual acts as something else. However, it denies that Impure Non-Identity Cases must be treated *either* like Pure Non-Identity Cases or as Pure Harming Cases (i.e., that negative wellbeing effects must be given a W or 2W weighting). Rather, Option 4 invites us to see these cases as somewhere in between – harder to justify than Pure Non-Identity Cases, but easier to justify than Pure Harming Cases. We could see how hard some negative wellbeing effect is to justify as a scale – with pure harms at one end, and pure impersonal losses at the other. And the considerations I canvassed above might help us place Impure Cases somewhere along this scale.

In closing this section, I would like to make clear an interesting entailment of accepting either Option 4, or a view under which we sometimes accept Option 2 and sometimes Option 3. This entailment concerns intervening agency. In these cases, intervening agency has the *opposite* significance from that which it usually has. Ordinarily, intervening agency, if it has any moral effect, makes harm *easier* to justify, than if one brings it about oneself. It isn't as bad to sell a gun to someone one knows might use it to unjustifiably harm someone as it is to directly risk unjustifiably harming someone. Here the intervening agency makes things *harder to justify*. And the stronger, more independent, that intervening agency is, the harder the act becomes to justify. If I merely foresee that my act may cause unrelated others to act in a certain way, I should view the negative wellbeing effects I cause as harms, or as

closer to harms in terms of justification. But if I order my subordinates to bring about negative wellbeing effects, or press a button bringing those about, I should view these as more like my own acts, and so I should treat them as impersonal losses, or more like impersonal losses in terms of justification.

## VI. Double Overdetermination Cases

In the introduction I briefly mentioned overdetermination of harm cases, such as the assassination of Julius Caesar. In introducing Impure Non-Identity cases I have also mentioned cases in which identity is overdetermined. In considering overdetermined harm cases, we kept identity constant. And in considering overdetermination of identity cases, we kept the individual negative wellbeing effects independent. But what if we put the two forms of overdetermination cases together?

That will give us cases such as this:

*Double Overdetermination:* If fewer than  $x$  acts of Type T occur, Person A (or Group A) will exist. If  $x$  or more acts of Type T occur, Person B (or Group B) will exist. If fewer than  $y$  acts of Type T occur, the person (or group) will have a higher level of wellbeing. If  $y$  or more acts of Type T occur, the person (or group) will have a lower level of wellbeing.

This sort of case has two important thresholds – a point ( $x$ ) which is an identity threshold, and a point ( $y$ ) which is a Negative Wellbeing Effect threshold point.

Here is a concrete example:

*Emitters:* Country X has a choice – it can either radically reduce its emissions, or not. If it does not radically reduce its emissions, two consequences are foreseeable. First, at some point, the level of emissions will cause the country's main source of income, a certain plant, will die out, and life will become less good (threshold A). Second, at some point, the level of emissions will cause flooding, which will cause a major migration, which will have a major effect on who exists in the next generation (threshold B). Each individual act of emitting contributes toward Country A reaching both of these thresholds.

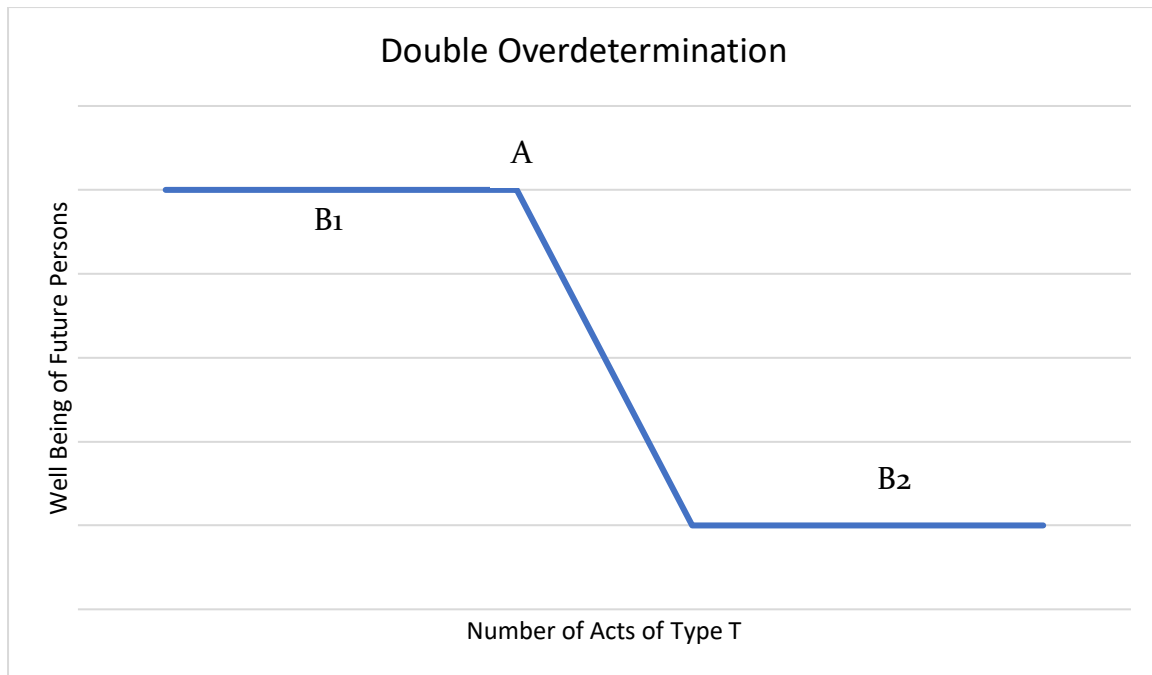
Double Overdetermination cases (cases in which both thresholds are passed by at least one act) are not Impure Non-Identity cases as I have defined them thus far. In Impure Non-Identity cases, the course of action is a non-identity case, and thus not harmful, but the individual acts are harmful. In Double Overdetermination cases, neither the course of action nor the individual acts are harmful.

Imagine, for example, that Country A does not reduce its emissions. The next generation, left badly off, seeks to complain to the older generation about their plight. Policy-makers can note that since, taken collectively, this is a non-identity case, and so the future generation has not been left any worse off. Individual emitters can note that, since the negative wellbeing effect was overdetermined, their acts were also not harmful – take any individual act of emitting away, and the same still would have occurred.

This is troubling. The response we'd ordinarily make, and that Parfit suggests we make, to one of the emitters in an overdetermined harm case would be to point out that it doesn't only matter what I have done, it matters what I have done *in conjunction with others*. Were Brutus to claim that he had not harmed Caesar, we would rightly reply that while that is true of his act *taken alone*, it is not true of the collection of acts of which his was a constituent member. But this reply is unavailable to us in Double Overdetermination cases. Had the *collection* of acts not taken place, the complainant wouldn't have existed.

Should Double Overdetermination cases be treated in the same way as Pure Non-Identity cases then? Do those caused to exist in a Double Overdetermination scenario lack any standing to complain? In what follows, I to try to construct an argument for those left badly off having complaints in at least some Double Overdetermination cases.

Consider the following graph.



The y axis represents the wellbeing of the group. The x axis represents the number of acts of Type T. While it is easy to think of the x axis as representing a temporally ordered series of acts, so that it in essence represents time, these acts could be simultaneous, or some of them might happen simultaneously.

At point A on the graph, we reach threshold A -- the Negative Wellbeing Effect threshold. Enough acts of Type T have been performed at that point that there is an irreversible decline in wellbeing. The identity tipping point, threshold B, could occur either before or after the negative effect tipping point, or exactly at the same point. It could occur at point B<sub>1</sub>, so that to the left of point B<sub>1</sub>, Group 1 will exist, and to the right of point 1, a different set of people, Group 2, will exist. Or it could occur at point B<sub>2</sub>, so that to the left of point B<sub>2</sub>, Group 1 will exist, and to the right of point B<sub>2</sub>, Group 2 will exist.

As we have seen, neither the individual acts nor the collection of acts harm in Double Overdetermination cases. But perhaps there is nevertheless some room for complaints in these cases. Consider, first, the ex post situation in the case in which the identity tipping point, threshold B, is at point B<sub>1</sub>. In this case, Group 2 may not have been harmed by any individual action or by the whole group of actions, but there was a possibility that they could have lived at the higher level. If there were enough Type T acts to get them past point B<sub>1</sub>, so that they existed, but not enough to get to threshold A, they would have had very good lives.

In cases in which agents are uncoordinated, Group B possibly has a complaint against each individual actor since each act raised the probability that they would exist but also raised the probability that they would live at the lower level, and not the higher level. Where agents are at least informed about each other's decisions, particularly weighty complaints might be pressed against those who acted after it became clear that the point B<sub>1</sub> threshold (the identity threshold) had been reached but not yet clear that the threshold A (the negative wellbeing effect threshold) would be reached. If the group acted collaboratively, a complaint could be pressed against the group, as there is a sub-set of acts (those beyond point B<sub>1</sub>) that collectively harm the Group B. The group could have stopped before they reached point A.

Consider, in contrast, the ex post situation in the case in which the identity tipping point is at point B<sub>2</sub>. In this case, there is no scope for complaint. This is more like a Pure Non-Identity Case, since, although what determines identity and what determines wellbeing are independent, the person or group who is now living at the

lower level could not have existed at the higher level. This is why it matters whether the identity threshold or the negative wellbeing effect threshold comes first.

Consider now the ex ante perspective. In the case where the identity tipping point is at point B<sub>2</sub>, then there is a risk of harming Group 1 – as things stand, they will exist, but there is a risk that by performing acts of Type T we will reach point A, but not point B<sub>2</sub>. There is a risk that Group 1 will exist *with a complaint*.

Here is an interesting question. Imagine we are at the point where we have passed the negative effect tipping point (point A) but have not yet reached point B<sub>2</sub>. Would we now have reason to perform more Type T acts so as to ensure that Group 2 exists? On the one hand, this looks like a straight choice between two groups who will live at the same level of wellbeing – there is nothing to choose between them. On the other, if Group 1 exists they will exist *possessing a complaint*, whereas Group B will not have a complaint. This raises an interesting question for those of us attracted to the Difference View – is it better to ensure that people exist with no complaint than that people exist with a complaint, even though they will be equally well off? I find the idea that we ought to push on with Type T acts, so as to ensure Group B exists counter-intuitive. It seems that Group A is being (further) punished for being wronged, or at least potentially wronged.

## VII. Concluding Remarks and Practical Applications

This is a work of abstract philosophy. I make no apologies for that. Even if the issues had no practical relevance, they are theoretically fascinating and important. But we



shouldn't be led by the abstract nature of the work, and the stylised examples, into thinking that these issues are of little or no practical importance. The topic we have here concerns courses of action that, together, are non-identity cases, but individually are not. I have focused here on cases in which *each and every individual act* is harmful, but the collection of acts is not. But Impure Non-Identity cases need not be so, well, pure. The important point is that some individual might owe their existence to a collection of acts, but not to each individual action within that collection. Here are two real world examples where the Impure Non-Identity Problem seems like it has relevance. The first is war. An individual act of war will have comparatively few identity-determining effects. Wars, on the other hand, clearly have far-ranging identity-determining effects. Wars can also make life very bad (or very good) for future generations.

The second is climate change. Both Michael Otsuka and John Broome describe climate change policy as an Impure Non-Identity Case (though they do not use that language). Otsuka notes that (as in Depletion) if we choose to reduce greenhouse gas emissions in the present, this will affect who later comes to exist. He writes:

If, therefore, we *collectively refrain* from taking the requisite greenhouse-gas-reducing measures, distant future people will not have any complaints of the stronger form that [according to the Difference View] only those whose existence is choice-independent can press.

The effects of the carbon footprint of a *single individual* might, however, give rise to such stronger complaints from those whose existence is independent of that individual's choices...[John] Broome maintains that it is 'extraordinarily

unlikely' that a typical individual's [lifetime] net emission of 800 tons of carbon dioxide will make *nobody* worse off than he would have been.<sup>13</sup>

Identifying Impure Non-Identity Cases, and discussing how our moral theories can and should handle them is of both theoretical and practical importance. I have tried to show here that Impure Non-Identity Cases exist, their structure, and the moral questions they pose.

## **Bibliography**

Broome, J., *Climate Matters: Ethics in a Warming World* (W.W. Norton, 2012).

Otsuka, M., 'How it makes a moral difference that one is worse off than one could have been,' in *Politics, Philosophy & Economics* 17 (2018): 192-215.

Parfit, D., *Reasons and Persons* (Clarendon Press, 1984).

Parfit, D., 'On What Matters: Volume Two' (Oxford University Press, 2011).

Parfit, D. 'Future People, the Non-Identity Problem, and Person-Affecting Principles' in *Philosophy & Public Affairs* 45 (2017): 118-157.

Roberts, M.A., 'The Nonidentity Problem,' in E.N. Zalta ed., *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition):

<<https://plato.stanford.edu/archives/win2015/entries/nonidentity-problem/>>.

Tomlin, P., 'Proportionality in War: Revising Revisionism' in *Ethics* (forthcoming).

---

<sup>13</sup> Otsuka, 'How it makes a moral difference.' pp. 209-210. Italics in original. Otsuka is drawing on Broome, *Climate Matters*, p. 64.