

A Non-Identity Dilemma for Person-Affecting Views

Elliott Thornley*

Abstract: Person-affecting views in population ethics state that (in cases where all else is equal) we're permitted but not required to create people who would enjoy good lives. In this paper, I present an argument against every possible variety of person-affecting view. The argument takes the form of a dilemma. Narrow person-affecting views must embrace at least one of three implausible verdicts in a case that I call 'Expanded Non-Identity.' Wide person-affecting views run into trouble in a case that I call 'Two-Shot Non-Identity.' One plausible practical upshot of my argument is as follows: we individuals and our governments should be doing more to reduce the risk of human extinction this century.

1. Introduction

My subject is person-affecting views in population ethics. As is custom, I begin with:

Narveson's Slogan

We are in favor of making people happy, but neutral about making happy people. (Narveson 1973, 80)

I'll take a sharpened version of the latter clause to define 'person-affecting views.' Person-affecting views are those views that imply the:

Deontic Principle of Neutrality

In cases where all else is equal, we're permitted but not required to create people who would enjoy good lives.¹

Intuitions about person-affecting views run both ways. These views seem less appealing when we note that future good lives could contain all the things that make our own lives valuable: joy, knowledge, achievement, loving relationships,

* Comments and questions welcome at elliott.thornley@philosophy.ox.ac.uk

¹ This principle is one half of the famous *Procreation Asymmetry* in population ethics (for which see McMahan 1981, 100; Holtug 2004; Roberts 2011a; 2011b; Chappell 2017; Frick 2017; 2020; Bader 2022; Thomas 2023; Thornley 2023; Francis, n.d.). The other half of the Procreation Asymmetry states that, in cases where all else is equal, we're required not to create people who would suffer bad lives.

and so on (Kavka 1978, 195–96). But person-affecting views seem more appealing when we instead note the following: if we decline to create a person who would enjoy a good life (in cases where all else is equal), then no existing person is worse off (Govier 1979, 111; Hare 2007, 498). From this perspective, declining to create a person looks like a victimless crime, which may lead us to believe that it is no crime at all.

Many of us feel the force of both of these intuitions. Other philosophers find one intuition compelling and the other unconvincing. Unfortunately, these other philosophers disagree about which intuition is which. Since these initial intuitions are a wash, we have to look at the arguments.

In this paper, I argue against person-affecting views. Arguments against these views have been given before, but most are tailored to the details of specific theories.² New variants of person-affecting views emerge unscathed. Other arguments employ premises that advocates of person-affecting views have proven happy to reject. The most famous argument against person-affecting views is a case in point. It begins with two claims common to such views:

Existence Anticomparativism

Existing can't be better for a person than not existing.

The Person-Affecting Restriction

An outcome can't be better than another unless it's better for some person.

These claims together imply that creating a person with a wonderful life is no better than creating a different person with a barely good life. It is not better for the person with the wonderful life (by Existence Anticomparativism) nor is it better for anyone else, and so it is not better *simpliciter* (by the Person-Affecting Restriction). That suggests (counterintuitively to many) that we are permitted to create a person with a barely good life rather than a different person with a wonderful life.

This *non-identity problem* (Parfit 1984, chap. 16) has long been considered the most serious objection to person-affecting views, but two developments cast doubt on its significance. The first is the growing number of philosophers who accept *narrow* person-affecting views' supposedly-unacceptable verdict that creating the person with the barely good life is permissible (Heyd 2009; Roberts 2011b; Boonin 2014; McDermott 2019; Mogensen 2019; Horton 2021; Podgorski

² See, for example, (Beckstead 2013, chap. 4; Ross 2015; Greaves 2017; Horton 2021; Thomas 2023; Thornley 2023; Arrhenius forthcoming, chap. 10).

2021; Spencer 2021). The second is the construction of *wide* person-affecting views which avoid the verdict (Hare 2007; Meacham 2012; Parfit 2017; Frick 2020).

In response to these developments, I present a dilemma for person-affecting views that builds on the non-identity problem. The first horn of this dilemma is a trilemma for narrow views. They must embrace at least one of three implausible verdicts in a case that I call ‘Expanded Non-Identity.’ The second horn of the dilemma is a sequential choice problem for wide views. A case that I call ‘Two-Shot Non-Identity’ reveals that these views make permissibility depend on factors that seem morally irrelevant: factors like whether we could have had a happier child many years earlier, or whether we carry out our choices by pulling two levers or one.

The premises of my argument are (I think) harder to reject than the premises of the original argument based on the non-identity problem. And my conclusion is that the Deontic Principle of Neutrality is false: in cases where all else is equal, we are required to create people who would enjoy good lives. My argument thus presents a challenge to all possible person-affecting views.

Now for one plausible practical upshot of my argument. There’s a risk that humanity goes extinct this century, and it’s widely agreed that the interests of existing people give us some reason to reduce this risk (Shulman and Thornley forthcoming). But if person-affecting views are false, then some *non*-person-affecting view in population ethics must be true, and many of these latter views imply that the prospect of happy future generations gives us additional reason to reduce the risk of human extinction this century (see, for example, Greaves 2017; Greaves and Ord 2017, sec. 4.4; Mogensen 2021, sec. 2). And since there could well be a lot of future generations enjoying very good lives, many of these views imply that this additional reason is strong (Tarsney and Thomas 2020; Greaves and MacAskill 2021). That in turn suggests that we individuals and our governments should be doing more to reduce the risk of human extinction this century.

2. Scene-setting

Recall that I’m defining ‘person-affecting views’ as those views that imply the:

Deontic Principle of Neutrality

In cases where all else is equal, we’re permitted but not required to create people who would enjoy good lives.

The Deontic Principle of Neutrality implies that we're permitted to choose either option in the following case:

Just Bobby

(1) Bobby 100

(2) —

Here option (1) is creating Bobby with a wonderful life, represented by a welfare level of 100. Option (2) is creating no one, represented by the '—'. The Deontic Principle of Neutrality implies that each of (1) and (2) is permissible.

I'll argue that the Deontic Principle of Neutrality is false, and hence that person-affecting views are false too. And my argument takes the form of a dilemma. To see the two horns, consider the following case:

One-Shot Non-Identity

(1) Amy 1

(2) Bobby 100

Here option (1) is creating Amy with a life that is just barely good at welfare level 1. Option (2) is creating Bobby with a wonderful life at welfare level 100. We can split person-affecting views into two classes based on their verdicts in cases like One-Shot Non-Identity: cases in which we must either create a person with a good life or create a different person with a better life. The first class is:

Narrow person-affecting views

Those person-affecting views that imply that we are permitted to create the worse-off person in cases like One-Shot Non-Identity.

And the second class is:

Wide person-affecting views

Those person-affecting views that imply that we are required to create the better-off person in cases like One-Shot Non-Identity.³

I'll argue that each class of person-affecting view faces a serious challenge. Narrow views imply a trilemma. Wide views have trouble with sequential choice.

³ The 'narrow' and 'wide' terminology comes from Parfit (1984, chap. 18). My use matches that of Thomas (2023, 490).

3. The trilemma for narrow views

The defining verdict of narrow views might seem implausible, but many philosophers have made peace with it.⁴ I now present a harder problem for narrow views. Consider:

Expanded Non-Identity

- (1) Amy 1
- (2) Bobby 100
- (3) Amy 10, Bobby 10

Here we've added a third option to One-Shot Non-Identity. The first two options are as before: create Amy with a barely good life at welfare level 1 or create Bobby with a wonderful life at welfare level 100. The new third option is to create both Amy and Bobby with mediocre lives at welfare level 10.

Narrow views imply that each of (1) and (2) is permissible when these are the only available options. What should they say when (3) is also an option? I'll argue that they must say at least one of three implausible things, so that narrow views face a trilemma.

3.1. Option (1) remains permissible

The first thing they could say is that option (1) – creating Amy with a barely good life at welfare level 1 – remains permissible when we move from One-Shot Non-Identity to Expanded Non-Identity. But that claim implies:

Permissible to Choose Dominated Options

There are option sets in which we're permitted to choose some option X even though there's some other available option Y such that (i) everyone in X is better off in Y , (ii) everyone who exists in Y but not X has a good life, and (iii) Y is perfectly equal.

That's because (1) is *dominated* by (3): (3) creates only people with good lives, it leads to perfect equality, and it's better than (1) for Amy: the only person who exists in (1).

Permissible to Choose Dominated Options seems implausible. I think it's so implausible that we have decisive reason to reject any view that implies it. But a possible response goes like this: although it would be implausible to claim that

⁴ See, for example, Heyd 2009; Roberts 2011b; Boonin 2014; McDermott 2019; Mogensen 2019; Horton 2021; Podgorski 2021; Spencer 2021.

(1) is permissible in a straight choice between (1) and (3), it is not so implausible to claim that (1) is permissible in Expanded Non-Identity where (2) is also an option. In cases where (2) is also an option, (3) *harms* Bobby because he's better off in (2).⁵ So although (3) dominates (1), it's permissible to choose (1) in Expanded Non-Identity, because (3) harms Bobby in Expanded Non-Identity.

I find this argument unconvincing. The fact that Bobby is better off in (2) gives us reason to choose (2) over (3). It doesn't give us reason to choose (1). After all, Bobby enjoys a good life in (3) and in (1) he doesn't exist at all.⁶ And even if it did give us some reason to choose (1), we'd have to weigh this reason against our reason *not* to choose (1): (1) is worse for Amy than (3) and it brings about much less welfare than each of (2) and (3). Imagine the conversation you might have with Amy after choosing (1):

Amy: Why did you harm me? Why did you make me worse off than I could have been?

You: Because to make you better off, I would have had to create another person with a good life.

Amy: What's wrong with that?

You: I had a third option: making this person's life even better.

Of course, choosing (3) might lead to a similar conversation with Bobby:

Bobby: Why did you harm me? Why did you make me worse off than I could have been?

But here your response is 'Because to make you better off, I would have had to not create Amy.' Arguably, this justification is at least a little more convincing. But if you're not convinced, then the right conclusion to draw is *not* that (1) and (3) are both permissible. It's that only (2) is permissible.⁷ So, I contend, the verdict that (1) is permissible in Expanded Non-Identity remains seriously implausible.

⁵ As the italics indicate, 'harm' is here used in a somewhat technical sense: your choice *harms* a person iff some other available choice would have left them better off. See Roberts (2011b, 337), Meacham (2012, 260), McDermott (2019, 437), and Thomas (2023, 473).

⁶ Supposing that we have reason to choose (1) over (3) leads to what Ross calls 'the Problem of Improvable-Life Avoidance': 'we have... reason to prefer outcomes in which a given person does not exist to outcomes in which this person exists and has an improvable life' (2015, 443).

⁷ This verdict is (I think) the correct one. But as we'll see in section 3.3, it leads to trouble when coupled up with any narrow view.

3.2. Option (3) is permissible

Here's something else that narrow views could say about Expanded Non-Identity: option (3) – creating Amy and Bobby with mediocre lives at welfare level 10 – is permissible. But that claim implies:

Permissible to Prioritise Creating Mediocre Lives:

There are option sets in which we're permitted to choose some option X even though there's some other available option Y such that X is much worse for the only person in Y and mediocre for the only person in X but not Y .

That's because (3) is mediocre for Amy and much worse than (2) for Bobby: Bobby's welfare level is 100 in (2) and 10 in (3). And we can imagine variations on Expanded Non-Identity in which Bobby's welfare level in (2) is arbitrarily high. The higher Bobby's welfare level in (2), the more implausible it is to claim that we're permitted to choose (3).

Here's another thing to consider. One objection often made to non-person-affecting views in population ethics is that these views will sometimes countenance making particular people worse off in order to create more people. But if a narrow person-affecting view says that (3) is permissible in Expanded Non-Identity, this view is *very* permissive about making particular people worse off in order to create more people. It permits us to drop Bobby's welfare level down from 100 to 10 in order to create Amy at welfare level 10. The narrow view in question is more permissive in this respect than even total utilitarianism. On total utilitarianism, choosing (3) is wrong because (2) leads to greater total welfare.

Here's another argument that choosing (3) is wrong: choosing (3) would certainly be wrong in a straight choice between (2) and (3), and adding (1) as an option can't make (3) permissible.⁸ Imagine the conversation you might have with Bobby after choosing (3):

Bobby: Why did you seriously harm me? Why did you make me much worse off than I could have been?

You: Because doing so allowed me to create Amy.

Bobby: That doesn't sound very person-affecting of you. Oh well, at least Amy's life must be very good.

⁸ In other work, I've called the negation of this last claim 'the Problem of Impairable-Life Acceptance' (Thornley 2023, sec. 6).

You: Oh, it's mediocre actually.

Bobby: What! How can it be permissible to seriously harm me in order to create Amy with a mediocre life?

You: Well, you see, I had another option: creating Amy with an even more mediocre life.

A possible response goes like this: the availability of (1) does make a difference, but not because Amy is worse off in (1) than in (3). The availability of (1) makes a difference because Bobby doesn't exist in (1). When (1) is available, Bobby's existence is contingent on your choice, so it's permissible to make Bobby's life much worse by choosing (3) rather than (2).

A superficial similarity with narrow views' verdict in One-Shot Non-Identity might make this claim seem defensible. But the case at hand isn't so much a non-identity problem as it is an *identity non-problem*. It's morally important not to make particular people worse off, even if you also have the option not to create them at all. So, I conclude, it's implausible for narrow views to say that (3) is permissible in Expanded Non-Identity.

3.3. Only option (2) is permissible

Now we can complete the trilemma for narrow views. If neither of (1) and (3) is permissible in Expanded Non-Identity, it must be that only (2) is permissible. But if only (2) is permissible, then narrow views imply:

Losers Can Dislodge Winners:

Adding some option X to an option set can make it wrong to choose a previously-permissible option Y , even though choosing X is itself wrong in the resulting option set.⁹

That's because narrow views imply that each of (1) and (2) is permissible in One-Shot Non-Identity. So if only (2) is permissible in Expanded Non-Identity, then adding (3) to our option set has made it wrong to choose (1) even though choosing (3) is itself wrong in Expanded Non-Identity.

That's a peculiar implication. It's a deontic version of an old anecdote about the philosopher Sidney Morgenbesser (Poundstone 2008, 50). Here's how that story goes. Morgenbesser is offered a choice between apple pie and blueberry pie, and he orders the apple. Shortly after, the waiter returns to say that cherry

⁹ This condition is the negation of Podgorski's (2021, 362) Losers Can't Dislodge Winners.

pie is also an option, to which Morgenbesser replies, ‘In that case, I’ll have the blueberry.’

That’s a strange pattern of preferences. The pattern is even stranger in our deontic case. Imagine instead that the waiter is offering Morgenbesser the options in Expanded Non-Identity.¹⁰ Initially the choice is between (1) and (2), and Morgenbesser permissibly opts for (1). Then the waiter returns to say that (3) is also an option, to which Morgenbesser replies, ‘In that case, I’m morally required to switch to (2).’ The upshot is that the waiter can force Morgenbesser’s hand by adding options that are wrong to choose in the resulting option set. And turning the case around, the waiter could expand Morgenbesser’s menu of permissible options by taking wrong options off the table. That seems implausible.

A possible response is as follows: the implication only sounds implausible at its current level of abstraction. A more concrete explanation makes it sound more plausible: choosing (3) is wrong because it’s much worse than (2) for Bobby, and adding (3) makes choosing (1) wrong because the availability of (3) means that (1) harms Amy.¹¹

I find this explanation unconvincing. It claims that it’s for Amy’s sake that Morgenbesser must switch from (1) to (2) when (3) is introduced. But (2) isn’t better for Amy. She enjoys a good life in (1), and in (2) she lives no life at all. Claiming that it’s for Amy’s sake that Morgenbesser must switch from (1) to (2) leads to what Horton calls ‘the problem of backfiring complaints’ (2021, 490): the implication that people’s complaints can backfire, making it wrong to create them even though they would enjoy good lives. In this case, it’s the introduction of (3) – an option that is better for Amy – that makes it wrong to create Amy at all.¹²

¹⁰ It’s a very unusual restaurant.

¹¹ Thanks to Jonas Hertel for this point.

¹² Here’s another issue. Those who claim that

adding (3) makes choosing (1) wrong because the availability of (3) means that (1) harms Amy.

seem committed to the following claims:

- (i) The fact that a person would be harmed in an outcome gives us reason *not* to choose that outcome.
- (ii) The fact that a person would enjoy a good life in an outcome gives us *no* reason *to* choose that outcome.

But these two claims together give rise to a problem. Meacham (2012, 281) calls the case ‘Asymmetric Creation.’ Horton (2021, 490) presents a similar case:

So it's implausible to claim that Morgenbesser is required to switch from (1) to (2) for Amy's sake. For whose sake is he required to switch? Not Bobby's. If Morgenbesser were required to choose (2) for Bobby's sake in Expanded Non-Identity, he'd presumably also be required to choose (2) for Bobby's sake in One-Shot Non-Identity, contrary to the defining verdict of narrow views.

Thus the narrow person-affecting views in question seem forced to say that we're required to choose (2) for no one's sake, and then it's debatable whether they're still worthy of the name: person-affecting views were supposed to avoid this air of impersonality. What's more, a question remains to be answered: if we're required to choose (2) for no one's sake in Expanded Non-Identity, why aren't we also required to choose (2) for no one's sake in One-Shot Non-Identity? A natural explanation of the antecedent is that (2) is the best option in Expanded Non-Identity: it leads to the most of what makes life good and the least of what makes life bad (or at least: it leads to the best balance of these things). But this explanation also suggests that (2) is the best option in One-Shot Non-Identity and thereby suggests that narrow views' defining verdict is false. The overarching lesson is that Losers Can Dislodge Winners remains implausible.

-
- (1*) —
 - (2*) Amy 99, Bobby 100
 - (3*) Amy 100, Bobby 99

Given claim (ii) above, Amy's and Bobby's good lives in (2*) and (3*) give us no reason to choose (2*) or (3*) over (1*). But given claim (i), the fact that Amy is harmed in (2*) – because she's better off in (3*) – gives us reason to choose (1*) over (2*). And the fact that Bobby is harmed in (3*) – because he's better off in (2*) – gives us reason to choose (1*) over (3*). So given claims (i) and (ii), we have reason to choose (1*) over (2*) and (3*), and we have no reason to choose (2*) or (3*) over (1*). Given a plausible principle linking reasons and obligations, we're then required to choose (1*). That verdict seems undesirable.

What's more, the verdict might seem even less desirable once we observe that cases of this kind are realistic and somewhat common. Arguably, anyone who finds themselves in the early stages of a pregnancy with twins is in the relevant situation. If you terminate the pregnancy early, Amy and Bobby will never exist. You'd thereby choose (1*). If you have the twins, you'll have to make some choice about how to divide your time, money, and attention between them. This choice will almost certainly have some effect on their welfare, making your choice of division akin to a choice between (2*) and (3*). And even a perfect 50-50 split of time, money, and attention will harm each twin in the relevant sense, because you could have made each twin better off by skewing the split their way.

So if we want to avoid the verdict that having twins is wrong, we should reject views that are committed to claims (i) and (ii). That in turn should make us wary of the claim at the start of this footnote: adding (3) makes choosing (1) wrong because the availability of (3) means that (1) harms Amy.

3.4. Summarising the trilemma

Now the trilemma for narrow person-affecting views is complete and I can summarise. If these views say that (1) is permissible in Expanded Non-Identity, they imply that it's Permissible to Choose Dominated Options. If they say that (3) is permissible, they imply that it's Permissible to Prioritise Creating Mediocre Lives. And if they say that only (2) is permissible, they imply Losers Can Dislodge Winners. Each of these implications is seriously implausible.¹³

This trilemma for narrow views is the first horn of the dilemma for person-affecting views as a whole. The other horn is a sequential choice problem for wide person-affecting views.

4. The sequential choice problem for wide views

Recall:

Wide person-affecting views

Those person-affecting views that imply that we are required to create the better-off person in cases like One-Shot Non-Identity.

Wide views imply that only (2) is permissible in One-Shot Non-Identity and so avoid the trilemma above: they can say that only (2) is permissible in Expanded Non-Identity without implying Losers Can Dislodge Winners. However, wide views have trouble with sequential choice. To see how, note the following.

¹³ The trilemma also has the virtue of being adjustable. We can make it harder to believe that (1) is permissible by increasing the welfare levels of Amy and Bobby in (3). For example, we could have option (3) be creating both Amy and Bobby with decent lives at welfare level 20 (rather than mediocre lives at welfare level 10). Or we can make it harder to believe that (3) is permissible by decreasing the welfare levels of Amy and Bobby in (3). For example, we could have option (3) be creating both Amy and Bobby with just-better-than-barely-good lives at welfare level 2.

We can also make it harder to believe that certain options are permissible by adding colour to the case. Our intuitions sometimes vary depending on whether welfare-decreases come about as a result of removing good things or as a result of adding bad things, and we can use this fact. For example, it will seem especially implausible to claim that (1) is permissible if Amy's life in (1) is exactly like her life in (3) except with enough suffering tacked on at the end to bring her welfare level down from 10 to 1. And it will seem especially implausible to claim that (3) is permissible if Bobby's life in (3) is exactly like his life in (1) except with enough suffering tacked on at the end to bring his welfare level down from 100 to 10.

So if you find yourself thinking that (1) is much more plausibly permissible than (3) or vice versa, you should adjust the trilemma in the aforementioned ways, so as to make (1) and (3) seem equally (im)plausibly permissible. You thereby make the trilemma maximally challenging for narrow person-affecting views.

Because wide views are a class of person-affecting view, they imply that each option is permissible in the following two cases:

Just Amy

- (1) Amy 1
- (2) —

And:

Just Bobby

- (1) Bobby 100
- (2) —

Now suppose that we face these choices one after the other in a case that we can call:

Two-Shot Non-Identity

We make a choice in Just Amy, lock that choice in, and then move on to Just Bobby.

Suppose that we choose to create Amy in Just Amy and then later decline to create Bobby in Just Bobby. In that case, we've done something with effects on Amy and Bobby equivalent to the effects of choosing (1) in:

One-Shot Non-Identity

- (1) Amy 1
- (2) Bobby 100

In each case, we've created Amy at welfare level 1 and declined to create Bobby at welfare level 100. Wide views imply that choosing (1) in One-Shot Non-Identity is wrong. So what should wide views say about creating Amy and then later declining to create Bobby in Two-Shot Non-Identity?¹⁴

Here are two things that wide views could say:

¹⁴ Spencer (2021, 3837–38) briefly considers a case of this kind. He argues from the Procreation Asymmetry, an agglomeration principle, and a variant of Sen's (2017, 63) Alpha to the conclusion that we should reject wide views and bite the bullet on the non-identity problem. His is a narrow view that is subject to my trilemma above.

Permissive wide views

We're *always* permitted to create the worse-off person and to later decline to create the better-off person in cases like Two-Shot Non-Identity.

Restrictive wide views

We're *never* permitted to create the worse-off person and to later decline to create the better-off person in cases like Two-Shot Non-Identity.

I'll take these two classes of view in turn. I'll then consider views that chart a course between them.

4.1. Permissive wide views

Here's a first problem for permissive wide views. Wide views were supposed to avoid counterintuitive verdicts in non-identity cases, but permissive wide views imply the counterintuitive verdict in something like Parfit's archetypal non-identity case, in which a prospective parent can have a worse-off child now or a better-off child later (Parfit 1984, 358). That prospective parent's predicament is more like Two-Shot Non-Identity than it is One-Shot Non-Identity and so permissive wide views seem to imply that having the worse-off child is permissible.¹⁵

Here's a more serious problem for permissive wide views: they make permissibility depend on factors that seem morally irrelevant. To see how, suppose that Amy's and Bobby's existence will be determined by the positions of two levers. By leaving the left lever up, we decline to create Amy. By pulling the left lever down, we create her at welfare level 1. By leaving the right lever up, we create Bobby at welfare level 100. By pulling the right lever down, we decline to create him. We first decide whether to create Amy, lock in that decision, and then decide whether to create Bobby.

¹⁵ Parfit's (1984, 358) presentation of the case is ambiguous. It doesn't make clear whether the prospective parent has the option to create both children. But we can imagine a version of Parfit's case in which this option is available. Suppose that, in this case, the parent first creates the worse-off child and then later declines to create the better-off child. Then permissive wide views imply that the parent did no wrong, but I expect many people's intuitions to demur.

Two-Shot Non-Identity with Levers



As it stands, the case is Two-Shot Non-Identity, and our permissive wide view implies that we are permitted to pull the left lever (thereby creating Amy) and then pull the right lever (thereby declining to create Bobby). But now suppose that someone lashes the two levers together so that our only options are pulling both levers or neither. Then our predicament is transformed into One-Shot Non-Identity and our permissive wide view implies that pulling both levers is wrong. That's a strange combination of verdicts. Pulling both levers should either be permissible in both cases or wrong in both cases. It shouldn't matter whether we can pull them one after the other.¹⁶

4.2. Restrictive wide views

So consider instead restrictive wide views, according to which we're never permitted to create the worse-off person and to then later decline to create the better-off person in cases like Two-Shot Non-Identity. Perhaps the latter choice is made wrong by the former, or perhaps each choice is permissible but the

¹⁶ Here's a possible reply, courtesy of Olle Risberg. What we're permitted to do depends on lever-lashing, but not because lever-lashing precludes pulling the levers one after the other. Instead, it's because lever-lashing removes the option to create both Amy and Bobby, and removes the option to create neither Amy nor Bobby. If we have the option to create both and the option to create neither, then creating just Amy is permissible. If we don't have the option to create both or the option to create neither, then creating just Amy is wrong.

This reply might have some promise, but it won't appeal to proponents of wide views. To see why, consider the following four-button case. By pressing button 1, we create just Amy with a barely good life. By pressing button 2, we create just Bobby with a wonderful life. By pressing button 3, we create both Amy and Bobby. By pressing button 4, we create neither Amy nor Bobby. The reply implies that it's permissible to create just Amy. That verdict doesn't contradict the letter of wide views (at least given my definition in this paper), but it certainly contradicts their spirit.

sequence is not.¹⁷ These views have implications that are unlikely to be welcomed by those inclined towards person-affecting views. They imply that a parent who previously chose to create Amy in *Just Amy* now *must* create Bobby in *Just Bobby* to avoid wrongdoing (even if Amy's creation occurred many decades in the past). That verdict runs counter to a major motivation for person-affecting views: permitting each option in cases where we can create someone who would enjoy a good life.

Another problem is one that restrictive wide views share with their permissive counterparts: making permissibility depend on factors that seem morally irrelevant. This time the factor is what one chose in the (possibly distant) past. To see the problem, suppose that a friend is considering having a child and comes to you for moral advice. Per restrictive wide views, you not only need to ask your friend the usual questions about the child's likely quality of life and how the child might affect existing people. You also need to ask your friend about their past procreative choices. If in the past your friend had a child with a worse life than this new child would have, your friend *must* have the new child to avoid wrongdoing. And now reversing the order of the cases: if in the past your friend declined to have a child with a better life than this new child would have, your friend *must not* have the new child. This latter implication seems to me especially counterintuitive. The new child's life could be wonderful, but if your friend previously declined to have a child with an even better life, your friend is not permitted to create them. Restrictive wide views imply that there are cases in which we are not even permitted to create a person who would enjoy a wonderful life.

4.3. Intermediate wide views


Given the defects of permissive and restrictive views, we might seek an intermediate wide view: a wide view that is sometimes permissive and sometimes restrictive. Perhaps (for example) wide views should say that there's something wrong with creating Amy and then later declining to create Bobby in *Two-Shot Non-Identity* if and only if you *foresee* at the time of creating Amy that you will later have the opportunity to create Bobby. Or perhaps our wide view should say that there's something wrong with creating Amy and then later declining to create Bobby if and only if you *intend* at the time of creating Amy to later decline to create Bobby.

¹⁷ This kind of sequence view has been considered in the literature on imprecise credences. See, for example, Weatherson (1998) and Elga (2010).

Views of this kind give more plausible verdicts in the previous cases – both the lever case and the enquiring friend case – but any exoneration is partial at best. The verdict in the friend case remains counterintuitive when we stipulate that your friend foresaw the choices that they would face. And although intentions are often relevant to questions of blameworthiness, I’m doubtful whether they are ever relevant to questions of permissibility. Certainly, it would be a surprising downside of wide views if they were committed to that controversial claim.¹⁸

But in any case, there’s a more serious objection waiting for all forms of intermediate wide view: like permissive and restrictive views, they make permissibility depend on factors that seem morally irrelevant. We can prove this by contradiction. Assume for *reductio* (i) any kind of wide person-affecting view, (ii) that permissibility doesn’t depend on lever-lashing, and (iii) that permissibility doesn’t depend on past choices. Consider again Two-Shot Non-Identity with Levers.

Two-Shot Non-Identity with Levers



¹⁸ See Rachels (1994), Thomson (1991, 293; 1999, 514–15), and Scanlon (2008, chap. 2), for arguments that intentions are irrelevant to permissibility.

Note also the following. For those who think that intentions are relevant to permissibility, it is the *moral character* of those intentions that matters: whether (for example) the intentions are good or bad, virtuous or vicious, pure or impure. But the intermediate wide view in question has permissibility depending on whether you intend to later create Bobby: an intention that by itself seems morally neutral. It seems especially implausible to claim that permissibility can depend on such morally neutral intentions. Thanks to Jakob Lohmar for this point.

When the levers are lashed together, the case is One-Shot Non-Identity and so wide person-affecting views imply that pulling both levers is wrong. If permissibility doesn't depend on lever-lashing, then it's also wrong to pull both levers when they aren't lashed together. So we can infer: if in the past we've pulled the first lever, it's wrong to pull the second lever. And if permissibility doesn't depend on past choices, then it's also wrong to pull the second lever in cases where we didn't previously pull the first lever. So pulling the second lever is wrong *simpliciter*. But pulling the second lever is declining to create Bobby, and so we get the conclusion that we're required to create Bobby, contrary to the Deontic Principle of Neutrality implied by wide person-affecting views. Since we began by assuming a wide person-affecting view, we've reached a contradiction and so we must reject at least one of our assumptions. Those who want to hold on to any kind of wide person-affecting view must admit that permissibility depends on lever-lashing or on past choices. Those who want to avoid this dependence must reject every kind of wide person-affecting view.

I expect few will want to argue that permissibility can depend on lever-lashing, but one might defend the verdict that permissibility can depend on past choices. Claiming that permissibility can never depend on past choices would rule out *resolute choice*: a popular approach to sequential decision-making. Per resolute choice, one should first decide on a permissible long-term plan and then choose options that accord with that plan, even if one needn't have chosen those options if one hadn't previously adopted the plan (McClennen 1990, sec. 9.6).

There's some appeal to the idea that choice should be forward-looking, and there are concomitant arguments against resolute choice (Hammond 1983, 183; Machina 1989, 1651–53; Gustafsson 2015, 1599; 2018, 602; 2022, chap. 7; Kowalczyk 2022, 15–16). But I need not make those arguments here. I need only argue that it's implausible for permissibility to depend on past choices *in cases like Two-Shot Non-Identity*. That claim is compatible with permissibility depending on past choices in other cases, like the cases that motivate resolute choice.¹⁹

And there are good reasons to think that permissibility doesn't depend on past choices in cases like Two-Shot Non-Identity. Here's one. If a wide view claims that permissibility depends on past choices, then we're required to create Bobby given that we've previously created Amy. But for whose sake? The two-shot nature of the case means that the answer can't be 'for Amy's sake': she's already created.²⁰ The only viable answers seem to be 'for Bobby's sake' and 'for

¹⁹ For example, in money pumps. See Gustafsson (2022).

²⁰ Ruling out this answer is a key innovation of Two-Shot Non-Identity. It's an important feature of the case, because philosophers' explanations of wide views' verdict in One-Shot Non-Identity

no one's sake,' and both answers spell trouble for person-affecting views. 'For Bobby's sake' sits uneasily with Existence Anticomparativism and 'for no one's sake' smells of impersonality. But more importantly, each answer applies just as well in the case in which you *didn't* previously create Amy, and so each answer suggests that you're required to create Bobby no matter what. That verdict is contrary to person-affecting views, and it implies that permissibility doesn't depend on past choices after all.

Here's a related point. Your past choice need not matter to Bobby: we can set up the case so that his life is the same no matter what you previously chose. And your present choice need matter only to Bobby: we can set up the case so that neither Amy nor anyone else is affected by the decision to create Bobby. Given all that, why should the permissibility of your present choice depend on your past choice? Any answer will reveal a concern for something besides Amy and Bobby: the people affected by your choices and the people whose existence is at stake. So having permissibility depend on past choices in Two-Shot Non-Identity implies an unseemly preoccupation with something that just doesn't matter to anyone whose interests are at stake.

This preoccupation is an undesirable feature of all forms of wide view. As is now clear, wide views sometimes remain undecided even when we know all the facts about who lives and how well. Their verdicts wait on the answers to questions that seem morally irrelevant: questions like 'Did you miss the opportunity to have a happier child many years earlier?' and 'Will you carry out your choice by pulling two levers or one?'

often appeal to the fate of Amy in some way. For example, Woodward (1986) argues that you're required to create Bobby in One-Shot Non-Identity because creating Amy instead would violate her rights. Shiffrin (1999) argues that you're required to create Bobby because creating Amy instead would illegitimately impose some harm on her in order to benefit her. Hare (2007) argues that you're required to create Bobby because creating Amy instead would be worse for your child *de dicto*. And Frick (2020) argues that you're required to create Bobby because creating Amy instead would violate the *Selection Requirement*:

In a choice between creating two possible persons, I have contrastive moral reason to create that person for whom I can better satisfy the moral standard that will obtain if I create that person. (Frick 2020, 79)

These points might explain wide views' verdict in One-Shot Non-Identity, but they cannot explain the verdict that we're required to create Bobby (given that we've already created Amy) in Two-Shot Non-Identity. The general reason is the same for each point: Amy is already created; her fate is already sealed; the only choice remaining is whether or not to create Bobby. So Woodward cannot appeal to Amy's rights to explain why we're required to create Bobby, nor can Shiffrin appeal to harms illegitimately imposed on Amy, nor can Hare appeal to what's worse for your child *de dicto*. And Frick's Selection Requirement doesn't apply, because the choice to create Bobby isn't 'a choice between creating two possible persons' (Frick 2020, 79).

5. Conclusion

My argument against person-affecting views is a garden of forking paths. The first fork is One-Shot Non-Identity: narrow views are those person-affecting views that permit us to create the worse-off person, and wide views are those person-affecting views that require us to create the better-off person.

The fork for narrow views is a trilemma in Expanded Non-Identity. These views imply Permissible to Choose Dominated Options, or Permissible to Prioritise Creating Mediocre Lives, or Losers Can Dislodge Winners.

The fork for wide views is a sequential choice problem. Setting aside objections to specific instances of wide views, these views make permissibility depend on facts that simply don't matter to anyone whose interests are at stake, like what you chose in the (perhaps distant) past and whether you make your choices by pulling two levers or one.

The conclusion is that the Deontic Principle of Neutrality is false: in cases where all else is equal, we're required to create people who would enjoy good lives. My argument thus presents a challenge to every possible variety of person-affecting view. As I note in the introduction, one plausible practical upshot is that we individuals and our governments should be doing more to reduce the risk of human extinction this century.²¹

6. References

- Arrhenius, Gustaf. forthcoming. *Population Ethics: The Challenge of Future Generations*. Oxford: Oxford University Press.
- Bader, Ralf. 2022. 'The Asymmetry'. In *Ethics and Existence: The Legacy of Derek Parfit*, edited by Jeff McMahan, Tim Campbell, James Goodrich, and Ketan Ramakrishnan, 15–37. Oxford: Oxford University Press.
- Beckstead, Nick. 2013. 'On the Overwhelming Importance of Shaping the Far Future'. PhD Thesis, Rutgers, New Jersey: Rutgers University. <http://dx.doi.org/doi:10.7282/T35M649T>.
- Boonin, David. 2014. *The Non-Identity Problem and the Ethics of Future People*. Oxford: Oxford University Press.
- Chappell, Richard Yetter. 2017. 'Rethinking the Asymmetry'. *Canadian Journal of Philosophy* 47 (2–3): 167–77.
- Elga, Adam. 2010. 'Subjective Probabilities Should Be Sharp'. *Philosopher's Imprint* 10 (5): 1–11.

²¹ For helpful comments and discussion, I thank Erik Carlson, Will Combs, Tomi Francis, Riley Harris, Jonas Hertel, Ina Jäntgen, David Lindqvist, Jakob Lohmar, Andreas Mogensen, Olle Risberg, Brad Saad, Rhys Southan, Luca Stroppa, Teru Thomas, Hayden Wilkinson, and audiences at LSE, St Andrews, and Uppsala.

- Francis, Tomi. n.d. 'In Favour of Making Happy People'. <https://drive.google.com/file/d/1XoVVH0kd6st50VspvLzZVZdA6uVKbk38/view>.
- Frick, Johann. 2017. 'On the Survival of Humanity'. *Canadian Journal of Philosophy* 47 (2–3): 344–67.
- . 2020. 'Conditional Reasons and the Procreation Asymmetry'. *Philosophical Perspectives* 34 (1): 53–87.
- Govier, Trudy. 1979. 'What Should We Do About Future People?' *American Philosophical Quarterly* 16 (2): 105–13.
- Greaves, Hilary. 2017. 'Population Axiology'. *Philosophy Compass* 12 (11).
- Greaves, Hilary, and William MacAskill. 2021. 'The Case for Strong Longtermism'. *GPI Working Paper* No. 5-2021. <https://globalprioritiesinstitute.org/hilary-greaves-william-macaskill-the-case-for-strong-longtermism/>.
- Greaves, Hilary, and Toby Ord. 2017. 'Moral Uncertainty About Population Axiology'. *Journal of Ethics and Social Philosophy* 12 (2): 135–67.
- Gustafsson, Johan E. 2015. 'Sequential Dominance and the Anti-Aggregation Principle'. *Philosophical Studies* 172 (6): 1593–1601. <https://doi.org/10.1007/s11098-014-0366-0>.
- . 2018. 'The Difference Principle Would Not Be Chosen Behind the Veil of Ignorance'. *Journal of Philosophy* 115 (11): 588–604. <https://doi.org/10.5840/jphil20181151134>.
- . 2022. *Money-Pump Arguments*. Elements in Decision Theory and Philosophy. Cambridge: Cambridge University Press.
- Hammond, Peter. 1983. 'Ex-Post Optimality as a Dynamically Consistent Objective for Collective Choice under Uncertainty'. In *Social Choice and Welfare*, edited by Prasanta K. Pattanaik and Maurice Salles, 145:175–205. Contributions to Economic Analysis. Elsevier. <https://doi.org/10.1016/B978-0-444-86487-1.50017-4>.
- Hare, Caspar. 2007. 'Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?' *Ethics* 117 (3): 498–523.
- Heyd, David. 2009. 'The Intractability of the Non-Identity Problem'. In *Harming Future Persons*, edited by Melinda A. Roberts and David T. Wasserman, 3–25. Dordrecht: Springer.
- Holtug, Nils. 2004. 'Person-Affecting Moralities'. In *The Repugnant Conclusion: Essays on Population Ethics*, edited by Torbjörn Tännsjö and Jesper Ryberg, 129–61. Dordrecht: Kluwer Academic Publishers.
- Horton, Joe. 2021. 'New and Improvable Lives'. *The Journal of Philosophy* 118 (9): 486–503.
- Kavka, Gregory. 1978. 'The Futurity Problem'. In *Obligations to Future Generations*, edited by R.I. Sikora and Brian Barry, 186–203. Philadelphia, PA.: Temple University Press.

- Kowalczyk, Kacper. 2022. 'People in Suitcases'. *Journal of Moral Philosophy* 20 (1–2): 3–30. <https://doi.org/10.1163/17455243-19030009>.
- Machina, Mark. 1989. 'Dynamic Consistency and Non-Expected Utility Models of Choice Under Risk'. *Journal of Economic Literature* 27 (February): 1622–68.
- McClellenn, Edward F. 1990. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511983979>.
- McDermott, Michael. 2019. 'Harms and Objections'. *Analysis* 79 (3): 436–48.
- McMahan, Jeff. 1981. 'Problems of Population Theory'. *Ethics* 92 (1): 96–127. <https://doi.org/10.1086/292301>.
- Meacham, Christopher J. G. 2012. 'Person-Affecting Views and Saturating Counterpart Relations'. *Philosophical Studies* 158 (2): 257–87.
- Mogensen, Andreas. 2019. 'Staking Our Future: Deontic Long-Termism and the Non-Identity Problem'. *GPI Working Paper* No. 9-2019. <https://globalprioritiesinstitute.org/andreas-mogensen-staking-our-future-deontic-long-termism-and-the-non-identity-problem/>.
- . 2021. 'Moral Demands and the Far Future'. *Philosophy and Phenomenological Research* 103 (3): 567–85. <https://doi.org/10.1111/phpr.12729>.
- Narveson, Jan. 1973. 'Moral Problems of Population'. *The Monist* 57 (1): 62–86.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- . 2017. 'Future People, the Non-Identity Problem, and Person-Affecting Principles'. *Philosophy & Public Affairs* 45 (2): 118–57. <https://doi.org/10.1111/papa.12088>.
- Podgorski, Abelard. 2021. 'Complaints and Tournament Population Ethics'. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12860>.
- Poundstone, William. 2008. *Gaming the Vote: Why Elections Aren't Fair (and What We Can Do About It)*. Hill and Wang.
- Rachels, James. 1994. 'More Impertinent Distinctions and a Defense of Active Euthanasia'. In *Killing and Letting Die*, edited by Bonnie Steinbock and Alastair Norcross, 2nd ed., 139–54. New York: Fordham University Press.
- Roberts, Melinda A. 2011a. 'An Asymmetry in the Ethics of Procreation'. *Philosophy Compass* 6 (11): 765–76.
- . 2011b. 'The Asymmetry: A Solution'. *Theoria* 77 (4): 333–67.
- Ross, Jacob. 2015. 'Rethinking the Person-Affecting Principle'. *Journal of Moral Philosophy* 12 (4): 428–61.
- Scanlon, T. M. 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Press.
- Sen, Amartya. 2017. *Collective Choice and Social Welfare*. Expanded Edition. London: Penguin.

- Shiffrin, Seana. 1999. 'Wrongful Life, Procreative Responsibility, and the Significance of Harm'. *Legal Theory* 5 (2): 117–48. <https://doi.org/10.1017/s1352325299052015>.
- Shulman, Carl, and Elliott Thornley. forthcoming. 'How Much Should Governments Pay to Prevent Catastrophes? Longtermism's Limited Role'. In *Essays on Longtermism*, edited by Jacob Barrett, Hilary Greaves, and David Thorstad. Oxford: Oxford University Press.
- Spencer, Jack. 2021. 'The Procreative Asymmetry and the Impossibility of Elusive Permission'. *Philosophical Studies* 178 (11): 3819–42.
- Tarsney, Christian J., and Teruji Thomas. 2020. 'Non-Additive Axiologies in Large Worlds'. <https://globalprioritiesinstitute.org/christian-tarsney-and-teruji-thomas-non-additive-axiologies-in-large-worlds/>.
- Thomas, Teruji. 2023. 'The Asymmetry, Uncertainty, and the Long Term'. *Philosophy and Phenomenological Research* 107 (2): 470–500.
- Thomson, Judith Jarvis. 1991. 'Self-Defense'. *Philosophy & Public Affairs* 20 (4): 283–310.
- Thomson, Judith Jarvis. 1999. 'Physician-Assisted Suicide: Two Moral Arguments'. *Ethics* 109 (3): 497–518. <https://doi.org/10.1086/233919>.
- Thornley, Elliott. 2023. 'The Procreation Asymmetry, Improvable-Life Avoidance, and Impairable-Life Acceptance'. *Analysis* 83 (3): 517–26.
- Weatherson, Brian. 1998. 'Decision Making with Imprecise Probabilities'. <https://brian.weatherson.org/vdt.pdf>.
- Woodward, James. 1986. 'The Non-Identity Problem'. *Ethics* 96 (4): 804–31. <https://doi.org/10.1086/292801>.