

Parental Obligations & the Non-Identity Problem

Jacob Isaac

University of British Columbia

Abstract: Since its proposal in 1984, Derek Parfit's 'Non-Identity Problem' has significantly influenced how social choice theorists understand existential harms and benefits. The 'problem' raises the question of whether parents act wrongly when they choose to create a child with a life barely worth living. It suggests that if the alternatives would have either led to a life not worth living or to non-existence, then the parents are not liable for moral criticism. This article challenges Parfit's premise by advocating for a Minimal Concern Clause that is generated through the parental relation. The Clause operates under the assumption that non-existence is not a valenced experiential state and is thus not comparable to any positive state of existence. The argument unfolds in two main steps: firstly, it disputes the premise in Same-Number Cases by proposing an alternative semantic distinction between objects of reference, particularly, the status of a merely possible person. Secondly, it contends that in Different-Number Cases, parents are obligated to exercise minimal prudence under conditions of uncertainty, and that the moral weight of the Minimal Concern Clause must be compounded to reflect the replicable nature of procreation.

1. Introduction

The non-identity problem is one of the most important and contested concepts in population ethics. Roughly, it argues that some actions which appear to be morally wrong, but which affect the identity of the conferee, are not wrong, since there is no one whom the action has made worse off. This kind of action can be seen in the following example:

Anais: Suppose Anais has always wanted a child and can conceive a child now or in three months. If she conceives now, her child (call them Child1) will be born with a congenital disorder which will significantly diminish their welfare and leave them with a life that is barely worth living. If Anais waits three months to conceive, she will have a different child (call them

Child2), who will be born with no health disorders and will live a very happy and fulfilling life. There are no downsides for Anais to wait to conceive, and she is aware of the reduced quality of life that her child will have if conceived now. Nevertheless, Anais is impatient and decides to conceive now.

Intuitively, Anais has done something wrong by knowingly conceiving a child with a serious health disorder, when she could have waited a short period of time and had a perfectly healthy child. But here we run into the issue of identity.

Because Child1 would be the result of a different spermatozoon fertilizing a different ovum than those which would result in Child2, the two possible children are different people. This follows from what Derek Parfit called the *Time-Dependence Claim*: “If any particular person had not been conceived when he was in fact conceived, it is in fact true that he would never have existed” (1984, pp. 351). Therefore, we cannot say that waiting to conceive would be better for Child1 since they would not exist; Child2 would exist instead. But, we also cannot say that waiting to conceive would be worse for Child1 since they would not exist and would thus lack the properties necessary to experience their own non-existence. We can refer to this as the No-Difference View.

Many theorists (Harman 2004, Hurka 2011, Temkin 2011, Frick 2022) accept the non-identity problem and argue that an action that brings about a life that is worth living, and which does not negatively affect any already existing person, is morally permissible. This includes creating lives that are *barely* worth living. Any inequalities or suffering present in such a life are thought to be outweighed by the existential benefits conferred on them. In other words, it makes no moral difference whether Anais conceives now or waits three months. I take this to be an unacceptable conclusion.

This paper intends to demonstrate that it is wrong to create a person whose welfare falls below a certain threshold, even if they have a life that is worth living and would not exist otherwise. My argument takes shape over two parts. First, I argue that in Same-Number Cases, those in which the same number of people will result regardless of which choice is made, it is wrong to cause a child to exist with low welfare when another child, with much higher welfare, could have existed in their place. To make this argument, I begin by challenging two dominant approaches in population ethics, before presenting a third approach that I believe is capable of overcoming their respective axiological issues, and which functions on a semantic distinction between objects of reference. In the second part of this article, I argue that in Different-Number Cases, relation-based obligations require that genesis choices be guided by minimal concern for conferees. To do so, I employ the approach established in part one to Jacob Nebel’s Intrapersonal Argument and present a new conception of minimal prudence under uncertainty. If I am successful, this article will show that it is always wrong to cause someone to exist with a life that is barely worth living.

2.1 Same-Number Cases

When we consider the non-identity problem, many of us feel as though we should be concerned with how these decisions impact the people who will be born. After all, is it not out of concern *for* Child1 that

1. This is a variation of Derek Parfit’s ‘14 Year-Old Girl’ case, which he presents in “Reasons and Persons” (Clarendon, D. Parfit, 1984), p. 358

that we think it wrong of Anais to conceive now? This thought can be summed up by what Parfit (1984) called:

The Narrow Person-affecting Principle: An outcome, O1, cannot be better (worse) than another outcome, O2, if there is no one for whom O1 is better (worse) than O2.

This view feels intuitively correct. Many of the things that we think of as morally good are good in nature of their being *good for* someone. We endorse the right to education and bodily autonomy because we think that they are good for the people who have them. Inversely, those things which we think are morally bad are so, because they are bad for someone. Consider sexual assault, murder, and slavery, for example. Yet, as we have just seen, the Narrow Person-affecting Principle leads to some counterintuitive results in cases that involve different identities. Specifically, in the non-identity problem, it suggests that conceiving Child1 cannot be worse than conceiving Child2 unless it is worse for someone, and we know that non-existence is not a state in which anyone can have a comparative experience. In cases that involve different people, then, the basic Narrow Person-affecting view offers unsatisfactory answers.

To further demonstrate the challenges that emerge with the narrow person-affecting principle, consider a variation of the non-identity problem presented earlier. In this case, the woman is *already* pregnant. She is told that if she moves ahead with the pregnancy her child will be born with a health disorder which will make life very difficult for them and ultimately less enjoyable. Nonetheless, the child will still have a life that is, on the whole, worth living. However, her doctor tells her that if she eats more vegetables during the first and second trimesters of her pregnancy, her child will not develop the health condition and will live a very happy and fulfilling life. The woman chooses *not* to eat more vegetables and her child is born with the serious health condition.

I expect that most of my readers will agree that the woman has done something wrong by knowingly bringing a child into existence with a seriously diminished quality of life when she could have easily done otherwise. The question then becomes: is this action morally worse than in Anais' case, where instead of having to eat more vegetables, she has to wait three months to conceive? I cannot see how the second case, involving the same identity across time, is worse than the original case which involves different identities. In both cases, the woman chooses to have a child with a diminished quality of life, despite there being no negative repercussions for her to either wait to conceive or eat more vegetables. Yet, the Narrow Person-affecting Principle would suggest that the former case is morally better than the latter since there is no one for whom Anais' action has made worse off.

Some may still think that the fact that the former case involves different possible identities *does* in fact make it morally better than the case involving an identity across time. But, consider a similar argument, one used by supporters of farming practices, often referred to as 'moral omnivorism'. One of the cornerstone arguments employed by moral omnivores is the *existence-dependence argument*, which claims that farming animals for food is not only morally permissible, but for some theorists, morally obligated, since those animals would not exist were it not for our farming practices. This idea can be traced as far back as 1896, when Leslie Stephen wrote:

"Of all the arguments for Vegetarianism, none is so weak as the argument from humanity. The pig has a stronger interest than anyone in the demand for bacon. If all the world were Jewish, there would be no pigs at all." (p. 236)

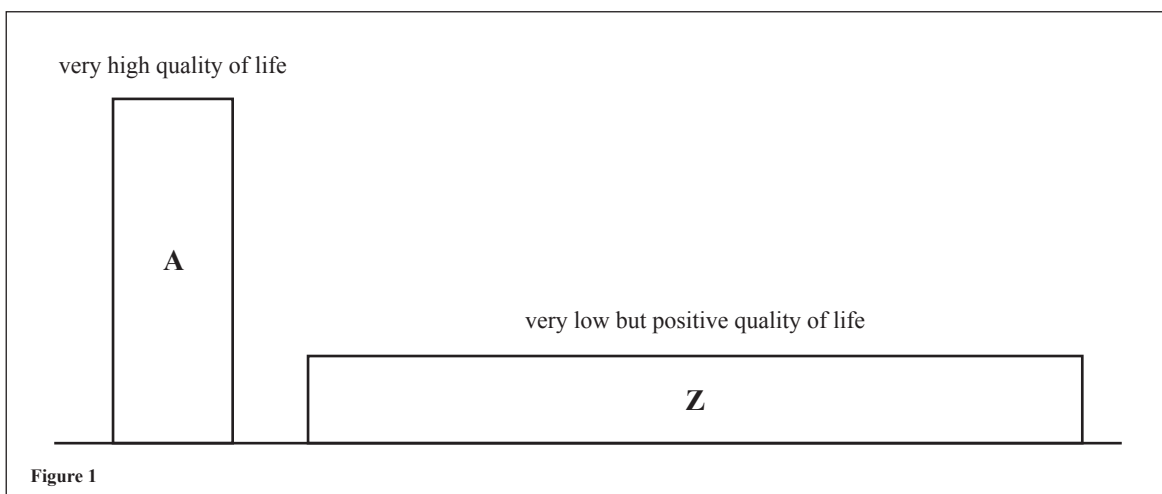
It is true that most of our current farming practices, such as factory farming, do not create animals with lives worth living, but let us, for a moment, humour moral omnivores and assume that they do. The argument now rests on the idea that the pigs would have no moral complaint (assuming they have the capacity to raise one), since they benefit from their existence, and their existence is dependent on their being raised for slaughter. The good they accrue from existing outweighs the bad they accrue from their eventual slaughter, so the argument goes. But framed this way, we could make the same argument about having a child for the *sole purpose* of abusing them for entertainment. If we only abuse them, say, once every three years, on the whole, this argument would go, they accrue more good from existing with a life that is, on the whole, worth living, than the bad they accrue from their triennial abuse.

I am unwilling to accept these arguments, and I hope my readers feel the same. By placing moral significance on the creation of new identities, and claiming that an act is not bad if there is no one who would exist in both states and who would be made worse off, the Narrow Person-affecting Principle justifies some rather concerning results.

Some who are hesitant to accept the Narrow Person-affecting Principle, instead, endorse an Impersonal Principle. On such a position, good (bad) things are not good (bad) for someone, but are good (bad) *simpliciter*. In the original formulation of the non-identity problem presented earlier, someone who maintains the Impersonal Principle would say that conceiving Child2 is better, not for Child2, but because it is better, all-things-considered, that there is more welfare in the world. The Impersonal Principle is a natural friend to total utilitarians, who claim that the genesis choices (those which will determine who exists in the future) we have a duty to make are those that bring about the most overall utility, expressed in this context as welfare. To find this overall utility, we simply add together the welfare of everyone existing in that world into one aggregate. This functions on what Parfit (1984) called:

The Impersonal Total Principle: If other things are equal, the best outcome is the one in which there would be the greatest quantity of happiness—the greatest net sum of happiness minus misery.

This means that a drop in the *quality* of lives in a given world can be counterbalanced by a larger increase in the *quantity* of lives. Yet, Parfit argued that such utilitarian thinking leads to the Repugnant Conclusion, shown in Figure 1.




The bars in Figure 1 represent two distinct populations, A and Z. The width of each bar represents the number of people that exist, with the height representing the welfare level of everyone in that population. Note that being above the x-axis, all of the lives in A and Z have positive welfare and thus are worth living (a life below the x-axis threshold would be sufficiently lacking welfare such that it would no longer be worth living).

According to Parfit, “For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better even though its members have lives that are barely worth living” (1984, pp. 388). The quality of life for those living in Z is significantly worse than in A but, due to Z’s massive population, there is a larger amalgam of welfare in Z. Consequently, although the people in A lead very good lives and the people in Z have lives only barely worth living, Z is nevertheless better than A according to the Impersonal Principle.

Some theorists (Sikora 1978, Broome 2004, Tännsjö 2020) have argued that the repugnant conclusion should not be considered so repugnant after all. Mackie (1985) argued that the repugnant conclusion simply points out the flaws in our initial intuitions, as we are quick to imagine the lives of those in Z as filled with immense suffering and devoid of any pleasures. But, the experiment makes clear that the people who exist in Z *do* have lives that are worth living. If they do in fact have lives that are worth living, perhaps Z should be seen as the population we ought to bring about. If this is correct, then the Impersonal Principle may present a tenable approach to questions of identity and welfare in population ethics. Yet, if taken to its logical end, the Impersonal Principle endorses the Very Repugnant Conclusion, as shown in Figure 2.

In the Very Repugnant Conclusion, we are presented with two worlds, A and Z. In A, there is a small population all living at a very high level of welfare. In Z, there are two groups of people, n and $n+$. The n group is larger than the total population of A, and all have lives that are not worth living, with extremely low, negative welfare. The $n+$ group has an immensely large population, all at very low levels of welfare, though still having lives that are worth living.



By the Impersonal Principle, $n+$ is so large that its sheer quantity counterbalances the negative welfare generated by the n group. On the whole, there are enough people in the $n+$ group that world Z is better than A.

I take Z in the Very Repugnant Conclusion to be worse than A. If we choose A, we will bring about a population of very well-off people to exist. If we choose Z, we will force a significantly larger population, those in the n group, to live lives of suffering. To imagine that by simply adding an immense number of people with lives that are still barely worth living, we can make this outcome better, I find unacceptable. This finds support from Jan Narveson's dictum that "morality favours making people happy, but is neutral to making happy people" (1973, pp. 73.) This has come to be known as the Procreation Asymmetry, or just 'The Asymmetry'.

A much more basic argument against the Impersonal Principle is that it is far too abstracted from what ought to drive our moral concerns. In Anais' case, if we were to say that it is worse if Child1 exists because it is better for *the world* if Child2 exists, we would be betraying our initial indignation at the thought of someone choosing to have a child with serious health defects when they could have chosen otherwise.

I hope that it is clear why neither a Narrow Person-affecting Principle nor an Impersonal Principle has the axiological strength to show why bringing people into existence with lives that are barely worth living is wrong. If we adopt an Impersonal Principle, we lose sight of the main concern for ethics, namely how people are impacted by our good (bad) decisions, and we come to endorse the repugnant conclusion. If, instead, we adopt a Narrow Person-affecting Principle, we place too much moral weight on transworld identities, and in doing so, consider any negative consequences of living at a low level of welfare as outweighed by the benefit of existence. The following section then, will present a third view, one that is capable of circumventing these challenges in Same-Number Cases.

2.2 The Wide Person-Affecting Principle

Overcoming the challenges associated with the Impersonal Principle, I believe, are relatively straightforward. We simply must fix our concern for welfare to the individuals for whom it accrues. But the Narrow Person-affecting Principle presents a larger hurdle, as our intuitions are often obscured by the confusion that emerges when we try to talk about possible people.

The language we use to answer questions involving different identities is important. Both the Narrow Person-Affecting Principle and the Impersonal Principle rely on the same semantic device when speaking about people: the *de re* quantifier. When we talk of a person in the *de re* sense, we are using a definite descriptor to refer to a proper noun. An available alternative descriptor that can be used when referring to someone, but which is precluded by the Narrow Person-Affecting and Impersonal Principles, is the *de dicto* quantifier. When we talk of a person in the *de dicto* sense, we are referring to a quasi-ambiguous descriptor, which relates to a non-binding variable. To demonstrate the difference, consider the following sentence: "Abril wants to marry an Italian woman". On a *de re* interpretation, we would say that Abril is already in love with a specific woman whom she wants to marry, and the woman happens to be Italian. 'Italian', in this case, is an adjective describing a proper noun (the specific woman). On a *de dicto* interpretation, we could say that Abril thinks that Italian women are particularly attractive and therefore wants the woman she ends up marrying to be Italian, though there is no specific Italian woman she currently wishes to marry.

This distinction is missed by the Narrow Person-affecting and Impersonal Principles with regard to transworld identities, as they only allow for semantic reference to specific, existing people *de re* (Hare, 2007). To see why this distinction matters, let us return to Anais' case.

Anais wishes to benefit her child. On a *de re* interpretation she does, since ‘her child’ is a definite proper noun referring to the person who was born; had she waited to conceive, this specific person would not have existed. Thus, on a *de re* interpretation, there is no moral issue with bringing Child1 into existence with a low level of welfare. The Narrow Person-affecting Principle requires only that she has concern for an actual person - her child *de re*, which is only possible after they exist. But ethical philosophy is concerned with decision-making; in the non-identity problem and the repugnant conclusion, we are tasked with choosing which outcome to bring about. If our decision-making is to involve consideration of how people *will fare* in the possible alternatives, we cannot limit our scope to actually existing people who can be captured by the *de re* quantifier.

A person-affecting principle that will allow for reference to possible people in the *de dicto* sense is thus necessary if we are to construct a welfare axiology that overcomes the challenges associated with trans-world identities. One such principle that I see as having intuitive merit is the:

The Wide Person-Affecting Principle: An outcome, O1, cannot be better (worse) than another outcome, O2, if there is no one for whom, were O1 to obtain, O1 would be better (worse) than O2 and no one for whom, were O2 to obtain, O2 would be worse (better) than O1. (Parfit, 1984).

This principle roughly states that outcome X would be worse than another outcome Y, if it benefits people less than outcome Y *would have* benefited people. The key here is that we can factor into our evaluation of an outcome, how much an alternative outcome would have benefited the people, *de dicto*, who would have resulted from that choice. In the non-identity problem, this would mean that we could say that bringing Child1 into existence is wrong because bringing Child2 into existence would have benefited Anais’ child (*de dicto*) more. Certainly, once Child1 exists, we should be solely concerned with their welfare. But a robust welfare axiology ought to guide our decision-making, and thus we need to be able to consider outcomes in *de dicto* terms prior to the action taking place.

Some scholars have rejected the wide person-affecting principle, on the grounds of moral actualism, which states that an outcome is morally wrong *iff* it makes *actually existing* people worse off. Inversely, moral actualism sees alternative outcomes that would affect merely possible people as morally irrelevant. David Boonin, in his 2014 book, *The Non-Identity Problem and the Ethics of Future People*, adopts this view, arguing that concern for merely possible people, accessed through *de dicto* semantic reference, should be ignored in our moral assessments of an outcome. To show this, he presents another variation of the non-identity problem, which involves Wilma, a prospective mother. In this case, Wilma is wanting to conceive a child and goes to visit her doctor for a check-up. Her doctor informs her that if she were to conceive now, her child would be born incurably blind and will have a seriously diminished quality of life, though, on the whole, it will still have a life that is worth living. However, if Wilma takes a pill for two months before conceiving, she will not have a blind child (Boonin, 2014, p. 2). If we apply the Wide Person-Affecting Principle and appeal to the quality of life that her child, *de dicto*, would have in both cases, we can point to the moral wrongness of Wilma’s choice to have her child be blind. However, Boonin presents a strong challenge to this thought.

If we are referring to Wilma’s child *de dicto*, we are not referring to any specific person, but the person who will eventually be known as Wilma’s child. Here, Boonin asks us to imagine that Wilma had instead decided to adopt a blind child. In this case, Wilma would be *choosing* for her child to be blind rather than sighted, because she could as easily have chosen to adopt a sighted child. He argues that making deontic

claims about choices involving de dicto people confuses our intuitions, since when we say that we have wronged someone, we are saying that they have been harmed in some relevant way. The Wide Person-Affecting Principle's appeal to de dicto people, he argues, allows us to divorce wronging someone from harming them, and thus fails to make serious deontic claims. He states that we must remain concerned with de re people, and thus the Narrow Person-Affecting Principle, since "It would clearly be absurd to think that Wilma's act of adopting a blind child wronged her child in virtue of the fact that it harmed her child in the de dicto sense. This would clearly seem to show that harming someone merely in the de dicto sense does not wrong them" (Boonin, 2014, p. 33).

While I think that this is a very strong argument, I do not think that it shows that wronging people de dicto is never morally significant. Consider, for example:

Biased doctor: Amanda is a doctor in an emergency room. On this particular day, the hospital is understaffed and unusually busy. There is a long line of patients all with serious injuries waiting for medical assistance. At the back of the line, Amanda sees one of her neighbours. They have a broken leg, a less serious injury than the other patients have. Since the hospital triages patients based on the severity of their injuries, Amanda's neighbour will have to wait a long time before being helped. Amanda is quite fond of her neighbour and calls them to the front of the line and helps them set their leg in a cast. Amanda's shift ends after she has helped her neighbour.

On a de re interpretation, Amanda has greatly benefitted her patients. She made sure that they got medical assistance much faster than they otherwise would have. But, I would argue that Amanda has done something wrong. She has harmed her patients (de dicto) by failing to triage in the correct order, instead choosing to help a friend rather than helping the patient most in need of medical assistance. It is clear that consideration for de dicto people can be morally significant. Further, Amanda's case shows that we often stand in obligation-generating relationships with de dicto people.

The suite of obligations generated through the parental relation includes *existential* obligations to de dicto children. Such obligations entail, among other things, not willingly giving your child an impairment which i) will significantly diminish their quality of life, and ii) is easily preventable at no cost to existing people (let us call this the Minimal Concern Clause). In the non-identity problem, Anais' choice to have a child, whether that be through conceiving now, in three months, or through adoption, generates in her these existential obligations towards the person who will eventually be her child. In choosing to conceive now, she fails to adhere to the Minimal Concern Clause, in a way that can be ignored in the adoption case. Despite what the Narrow Person-Affecting Principle suggests, simply bestowing a life that is worth living does not release you of all parental-based existential obligations.

Stated differently, it may be the case that the person (P1) who stands in relation to a parent when one option (O1) is chosen, is a different person than who stands in relation to that parent (P2) when the alternative option (O2) is chosen. In cases where the choice between O1 and O2 involves transworld identities and thus affects who will receive the existential benefits, that O1 is better for P1 than O2, says very little about the parent's adherence to their relation-based obligations in the case that they choose O1. However, the fact that Anais fails to respect the Minimal Concern Clause speaks directly to her lack of adherence to her existential obligations. In short, the obligation, generated through the parental relation, to ensure your child experiences a high welfare level is much more acute than the obligation to bestow a


life that is worth living.

It is true that all lives are a mixed bag of good and bad experiences; no one has ever lived a perfect life. What the Minimal Concern Clause allows us to determine, however, is whether the welfare level of the child chosen to be brought into existence is significantly diminished as a result of a parent's failure to fulfill their minimal obligations, or as the result of the natural challenges of living.

This thought pairs nicely with the Wide Person-Affecting Principle, since we can compare how much our choice to have a child with seriously diminished welfare would fulfil our parental obligations compared to our choice to have a perfectly well-off child. Through this principle, which allows for a consideration of *de dicto* people, we can say that Anais acts wrongly when choosing to have Child1, who will have a seriously diminished quality of life, rather than choosing to have Child2 with a very happy and fulfilling life. While she benefits her child *de re*, whichever child she chooses to conceive, she wrongs her child *de dicto* by choosing to conceive Child1.

2.3 A Challenge to the Wide Person-Affecting Principle

I said that my argument would take place in two steps. In the first section, I have demonstrated that the existential benefits of existing with a life that is barely worth living do not outweigh the wrongness of the decision to have such a child, in Same-Number Cases. This, I believe, has allowed us to identify, through the Wide Person-Affecting Principle, a solution to the non-identity problem, and a morally significant reason for Anais to conceive Child2. We may feel confident that the same principles are at play in *all* cases involving transworld identities. However, some theorists believe that we should discount relation-based existential obligations like the Minimal Concern Clause in Different-Number Cases. This is because the choice would no longer be between choosing to have a person with low welfare or a person with high welfare (both of whom could be captured by the *de dicto* reference 'my child'), but between bringing about someone with low welfare and not bringing about anyone at all. Out of sheer concern for the possible existential benefits conferred through a life that is worth living, they argue, we must discharge more demanding parental-based obligations such as the Minimal Concern Clause. Such a case is best typified by the Mere Addition Paradox, shown in Figure 3.



In the Mere Addition Paradox, there are three different possible worlds, A, B, and A+. Were we to bring about A, everyone living in that world would have a very high quality of life. World B is twice as large as that of A, but with a significantly lower welfare level. For anyone who is unconvinced by the

Impersonal Principle, and maintains that an increase in the quantity of lives *does not* counteract a decrease in the quality of lives, A will be better than B.

Half of A+'s population consists of the same people who existed in A, and their welfare has remained the same. The second half of A+'s population are contingent people, all with lives that are barely worth living. As the name suggests, Parfit (1984) wants us to consider the new group of people in A+ as being merely added to A. Since the population sizes of B and A+ are the same, and B has a higher average of welfare and fulfills the moral ideal of equality, B is better than A+.

The paradox emerges, however, when we compare A and A+. By transitivity of the better-than relation (which many social choice theorists see as an indispensable axiom), A should be better than A+. But the worst-off group in A+ have lives that are worth living and would otherwise not exist if we were to bring about A. Parfit argues that the contingent people in A+ have not been harmed by being brought into existence, since they have lives that are worth living, and if we were to choose A, they would not exist (1984, p. 420). Moreover, the contingent people in A+ were merely added, and as such their presence does not affect the better-off in A+ who existed in A. This logic, Parfit termed:

The Mere Addition Principle: Merely adding a group of new people, all of whom have lives that are worth living, does not make an outcome worse.
(1984, pp. 439).

The Mere Addition Principle is widely accepted by those who endorse the Narrow Person-affecting Principle. Since we ought to be concerned with how people, *de re*, fare, we should be satisfied with an outcome where they exist with lives that are worth living. On this view then, A+ is better than A. Likewise, because there are more people existing with lives worth living in A+, the Impersonal Principle also considers it better than A. This violates the transitivity of the better-than relation, since $A > B > A+ > A$.

The Mere Addition Paradox presents a challenge to our conception of parental-based obligations towards children *de dicto*. Perhaps it is correct to think that in Different-Number Cases where the only alternative to a life that is barely worth living is no life at all, we ought to discount the Minimal Concern Clause and place more moral weight on the obligation to bestow a life that is worth living.

The remainder of this article, then, will apply the Wide Person-Affecting Principle to Different-Number Cases, to show why this remains an unacceptable conclusion. It will do so by resolving Jacob Nebel's Intrapersonal Argument using the Minimal Concern Clause. Doing so, we will show that we ought not to bring about the worst-off in the Mere Addition Paradox.

3.1 Different-Number Cases

In *An Intrapersonal Addition Paradox*, Nebel (2019) presents a novel argument for the repugnant conclusion. Rather than reaching the repugnant conclusion through interpersonal principles, Nebel relies on intrapersonal principles. He attempts to demonstrate that a proper consideration for what is better *for individuals*, will also lead us to believe that Z is better than A in the repugnant conclusion. His argument rests on what he calls:

The Intrapersonal Repugnant Conclusion: For any person S, there exists some probability *p* such that any prospect in which S would have a

wonderful life with probability p or less, and would otherwise never exist, is worse for S than certainty of a life that is barely worth living. (2019, pp. 314).

For Nebel, a prospect that guarantees a life that is barely worth living is *always* better than a prospect that presents a less-than-perfect chance at perfect welfare. This means we ought to discount demanding existential obligations generated through the parental relation down to zero. He believes that this principle, functioning at the individual level can be extrapolated to show that it functions at the interpersonal level as well. He makes his argument in two steps.

First, he presents the case of a couple that is planning to conceive a child by injecting a single spermatozoon into a single ovum. He states that only one person (Sally) will be conceived through this process. Sally's parents have three options for the injection: A, Z, and A+ (shown in Table 1). If the parents choose A, Sally will have a very happy life at welfare level a if state 1 obtains, but will not exist if state 2 obtains. Option Z will give Sally a low quality of life at welfare level z . This will be the case if either state 1 or state 2 obtains. With option A+, if state 1 obtains, Sally will exist with a very high welfare level $a+$ (slightly above a). However, if state 2 obtains, Sally will have a life that is barely worth living at welfare level $z-$ (slightly below z).

	State 1 (p)	State 2 ($1-p$)
A	a	
A+	$a+$	$z-$
Z	z	z

Table 1

For Nebel, option A+ is better *for Sally* than option A, because it guarantees that she will exist with a life that is worth living, and if state 1 obtains, her welfare would be higher than it would be in A. Further, he thinks that Z would be better than A+, for some arbitrarily small p . By transitivity, then, he believes that Z will be better than A for Sally. But, here we should continue our discussion of relation-based obligations to de dicto people, by returning to the non-identity problem.

We have established that Anais acts wrongly by choosing to conceive Child1, giving them a life that is just barely worth living, rather than waiting three months and giving Child2 a much better life. It seems unacceptable to suggest that Anais would act wrongly if she were to choose to not have the child at all, rather than conceiving now. We know that nobody who does not exist can be harmed, since they lack the capacity to experience their non-existence. Further, we know that Child1 would have an incredibly hard life with significantly diminished welfare. If Anais conceived now, her child would have a moral complaint: their mother willingly gave them a very challenging life, one just barely worth living. But if Anais did not conceive at all, there would be nobody to raise such a complaint. It does not appear that simply by providing a life that is barely worth living, a parent fulfills even a minimal duty.

Notions of consent also play a significant role in assessing Anais' decision to have the child now. As mentioned earlier, existence is a mixed bag of good and bad, and just as no merely possible person

pursues existential benefits, neither do they consent to the bad that they will experience if they are conceived. In short, nobody chooses to be born. But would-be parents, at least in Nebel's experiment, have a rough awareness of the degree to which their child will be harmed by being brought into existence. If they choose to have the child, they consent, on behalf of their child, to the harms that they will experience.

Nebel seems to think that, if a life is, on the whole, worth living, parents are not liable for the harms that accrue in their child's life. This is akin to Joel Feinberg's argument that a rescuer cannot be held responsible for the injuries they caused to an endangered person. According to Feinberg, "the rescuer-defendant did not cause a condition that was harmful on balance, offset as it was by the overriding benefit of rescue. . . . [H]e cannot be said, therefore, to have harmed the [rescued person] (in the relevant full sense) at all." (1984, pp. 66). Imagine that, in rescuing the person (let us call them Ralph), the rescuer (Sam) breaks Ralph's leg. Both Nebel and Feinberg assume that a relevant, full sense of harm is one where someone acts against the greater interests of another. This creates a clear dichotomy between harms and benefits; benefits ultimately advance someone towards fulfilling their interests, and harms move them further away. Feinberg assumes that Ralph's ultimate interest is being saved from death, and though having his leg broken causes him pain, the cause of the broken leg (being rescued), contributes to the fulfilment of his interest. Likewise, Nebel assumes that existing with a life that is worth living is the ultimate interest of any rational person and that their parents have not harmed them in any relevant sense, by consenting to the significant harms associated with a life that is barely worth living.

Notice, however, that this conception of harms and benefits is dependent on a comparative starting point. If I were to break the leg of someone who is having a wonderful day with their spouse, I would be said to have harmed them. It is only because Ralph's starting position was so bad that their broken leg is not considered a harm. We know that non-existence cannot be good or bad for a person. Nor is it a state where someone can 'rationally pursue' a life. A life that is barely worth living is only non-comparatively good for the person who already has it, because at that point, they would not want to stop existing. They have a vested interest in its continuation, but knowledge of this interest cannot take lexical superiority over the decision to forcibly bestow it.

Nebel would have us consider *any* positive welfare, regardless of how much associated harm is experienced, benefits someone. But, as I have argued above, neutrality should not be our benchmark. The ultimate interest of any rational person, I believe, is not a marginally positive life, but a flourishing life. Applying our Wide Person-Affecting Principle, we can say that by consenting, on behalf of their child (de dicto), to the frustration of a flourishing life, which they could have guaranteed by picking A, Sally's parents harm her in a relevant sense. If they pick A, and State 2 obtains, they can try again. In intrapersonal cases, then, the Minimal Concern Clause should not be discounted, but rather compounded by each possibility to bring about your child (de dicto). Consenting on behalf of a child, who is not being harmed by their non-existence, to live a life that is incapable of flourishing is therefore wrong, full stop. I see no merit to Nebel's Intrapersonal Argument.

3.2 Obligations to Benefit

We have demonstrated that it would be wrong for Sally's parents to consent to Sally existing at welfare levels z and z^- . But a critic of the earlier section may say "If it is flourishing we ought to care about, have we not done something wrong by choosing for Sally to live at welfare level a rather than a^+ ?" This, I feel, has some intuitive appeal.

It is widely accepted in bio- and population ethics literature that parents *do* have some duties to benefit their child beyond a certain threshold. This has already been briefly shown through my arguments that we ought to benefit our dependents in ways beyond providing them with a life that is just worth living. Further, the obligations generated through the parental relationship are not discharged simply by providing shelter, food, and water. Seemingly ubiquitous in recent ethical works, Martha Nussbaum's *Capabilities Approach* argues that people have a moral right to ten basic functions that contribute to human flourishing, including affiliation, play, and bodily integrity (2013). All of these surpass the threshold of simply providing a child with a life worth living and subsequently keeping them alive (we will call this the existential threshold). We can therefore confidently assert that a parent who fails to perform their moral duty of benefitting their child beyond their existential threshold, and cannot provide sufficient reasons for doing so, is liable to moral criticism.

I would argue that the probability of state 2 obtaining (p being any arbitrarily small number) in the Intrapersonal Argument, is a sufficient reason for Sally's parents to not benefit Sally beyond welfare level a . If it would be more morally wrong for Sally's parents to consent to Sally living at welfare level z^- than to not conceive Sally at all, and they are aware that if they choose option A^+ , Sally will most likely exist at z^- , out of minimal prudence for the welfare of their child *de dicto*, they ought to choose option A . This of course flips what Nebel (2019, pp. 316) calls:

Minimal Prudence: For any individual S and very high welfare level x , there are some mediocre welfare levels y and y^- (where $y > y^-$) and some probability p such that some prospect in which S is certain to exist at level y is better for S than any prospect in which S might, with any probability less than or equal to p , exist at level x , and would otherwise exist at level y^- .

Here, we should instead consider Minimal Prudence as that which guarantees that someone will either be benefited beyond their existential threshold or otherwise will not be harmed at all. If state 1 obtains, both A and A^+ will ensure that Sally is benefitted beyond her existential threshold. However, if state 2 obtains (which is far more likely), A^+ will violate the Minimal Concern Clause and cause Sally harm, whereas A will not cause Sally to be harmed, since she would not exist. Sally's parents, therefore, have sufficient reason to not benefit Sally beyond welfare level a , and thus are not liable to moral criticisms. If they choose option A , they will be fulfilling their relational obligations; if they choose option A^+ , there is a very high likelihood that they will not. The minor increase from welfare level a to a^+ does not warrant such a risk.

3.3 Intra to Interpersonal Repugnance

I mentioned earlier that Nebel's argument happens in two steps. His first is establishing the Intrapersonal Argument, that Z is better than A for some arbitrarily small p . Having rejected this argument, I believe, we are in a position to reject the inductive step from intrapersonal to interpersonal cases, and therefore be able to solve for the Mere Addition Paradox. Nebel constructs his move from the intrapersonal to interpersonal repugnant argument on a set of social choice ideals. I will briefly touch on the first two of these ideals, both of which are also endorsed by our Wide Person-Affecting Principle, and demonstrate how his third negates his inductive step.

First, Nebel asks us to accept the *Same-Number Equality Claim*, which states that "Any two outcomes

containing the same number of people, all at the same level of well-being are equally good” (p. 318). This claim is compatible with our Wide Person-Affecting Principle, taking into account the interests of de dicto people. We affirmed this principle in our discussion of the non-identity problem.

Second, Nebel says that we must adhere to the *Stochastic Indifference for Equal Risk*, which states “For any egalitarian prospects X and Y, if every possible outcome of X and every possible outcome of Y are equally good, then X and Y are equally good” (p. 319). Egalitarianism, in this principle, refers to equal levels of welfare for all possible people, and an equal probability that any state would obtain. According to Nebel, “Rationality requires us to be indifferent between prospects that guarantee equally good outcomes - at least, when there is no risk of unfairness” (p. 319). Again, this principle squares nicely with our Wide Person-Affecting Principle, and I do not feel it needs further consideration.

Third, Nebel’s inductive step to interpersonal repugnance relies on a modified version of the classic Pareto Principle, which he calls *Weak Pareto for Equal Risk*, and which states “For any egalitarian prospects X and Y, if X is better than Y for each person who might exist in either prospect, then X is better than Y” (pp. 320). Nebel assumes this principle as a minimal condition of benevolence under certainty. Here, however, our Wide Person-Affecting Principle takes issue. Consider Table 2.

	State 1 (1/3)			State 2 (1/3)			State 3 (1/3)		
	Bob	Cat	Dan	Bob	Cat	Dan	Bob	Cat	Dan
A	a			a					a
Z	z	z	z	z	z	z	z	z	z

Table 2

Nebel states that “[w]e ought to prefer prospects that are better for everyone - at least, when there is no risk of unfairness” (p. 321). He takes this principle to mean that we ought to prefer prospects that are better for everyone *in the same obtaining state*. But, our Wide Person-Affecting Principle satisfies Weak Pareto for Equal Risk differently. All three states are equiprobable, and in outcome A+, each state gives Bob, Cat, and Dan an equal chance of existing at welfare level *a*.

We know that it would be wrong to bring one person into existence at welfare level *z*, since they could not possibly surpass their existential threshold. If we choose option A+, we guarantee that someone exists with a flourishing life; if we choose option Z, we guarantee that the Minimal Concern Clause is violated, constituting genuine harm in the lives of Bob, Cat, and Dan.

Notice, however, in Table 2 (which is taken directly from Nebel’s argument) and in my previous sentence, that Nebel presents this case in a way that intentionally confuses our intuitions: he gives the possible children names. Naming possible people, we are forced to refer to them as proper nouns, something that is only possible on a de re interpretation. And yet, at the time of decision-making, they are all merely possible children. Since we know that parental obligations extend to children de dicto, we must reject this intuition that they currently exist as Bob, Cat, and Dan. At the time of decision-making, we must solely be concerned with the welfare of the child(ren) de dicto who will exist. And we know, through our rejection of the Impersonal Principle, that more lives do not counterbalance worse lives. For these

reasons, we ought to prefer A+. This rejection of Nebel's interpretation of the Weak Pareto for Equal Risk stops his inductive step.

4. Conclusion

All of this is very technical, and it may appear that we have moved away from our initial question. But what we have shown is that in Nebel's first step, we can say that it is wrong to bring about a child who cannot pass their existential threshold, since the existential obligations generated through the parental relation entail fulfilling the Minimal Concern Clause. The Wide Person-Affecting Principle allows us to have concern for the child *de dicto*, who will come to exist, and any act which does not result in a child is not bad for that possible person, since they do not exist. Thus, we ought to prefer, for the sake of the child (*de dicto*) that they come to exist with a flourishing life.

In the second part of our argument, we blocked Nebel's inductive step by simply showing that a proper interpretation of Weak Pareto for Equal Risk does not entail that we should prefer more people existing with lives that are worth living, rather than one person with a flourishing life.

This can be easily extrapolated to the Mere Addition Paradox. It is wrong to consent on behalf of possible people to a life that cannot possibly provide benefits beyond their existential threshold. While these people will have a vested interest in the continuation of their life once they have it, the bad-making feature of the Mere Addition Paradox's A+ is the moral wrongness of the decision to bring about the worst-off. They would not have any interest in this meager life if they did not exist, and they would not be harmed in any relevant way, were they to not exist.

This argument is only possible through the application of a Wide Person-Affecting Principle which allows for concern for people *de dicto* prior to their creation. Our principle has shown that in both Same-Number Cases and Different-Number Cases, it is wrong to create lives that are barely worth living.

REFERENCES

1. Boonin, D. (2014). *Non-identity problem and the ethics of future people*. OXFORD UNIV Press.
2. Broome, J. (2004). *Weighing lives*. *Weighing Lives*, 1–19. <https://doi.org/10.1093/01992436x.003.0001>
3. Frick, J. (2022). *Context-dependent betterness and the mere addition paradox*. *Ethics and Existence*, 232–263. <https://doi.org/10.1093/oso/9780192894250.003.0010>
4. Hare, C. (2007). *Voices from Another world: Must we respect the interests of people who do not, and will never, exist?* *Ethics*, 117(3), 498–523. <https://doi.org/10.1086/512172>
5. Harman, E. (2004). *Can we harm and benefit in creating?* *Philosophical Perspectives*, 18(1), 89–113. <https://doi.org/10.1111/j.1520-8583.2004.00022.x>

6. Hurka, T. (2011). *Value and population size*. *Drawing Morals*, 25–36. <https://doi.org/10.1093/acprof:osobl/9780199743094.003.0002>
7. Mackie, J. L., 1985, “Parfit’s Population Paradox”, in J. Mackie & P. Mackie (eds.) *Persons and Values*, Oxford: Oxford University Press.
8. Narveson, J. (1973). *Moral problems of population*. *The Monist*, 57(1), 62–86.
9. Nussbaum, M. C. (2013). *Creating capabilities: The Human Development Approach*. Belknap Press of Harvard University Press.
10. Parfit, D. (1984). *Reasons and persons*. Clarendon.
11. Sikora, R. I. (1978). *Foundations without certainty*. *Canadian Journal of Philosophy*, 8(2), 227–245. <https://doi.org/10.1080/00455091.1978.10717048>
12. Stephen, L. (1896). *Social Rights and Duties*. Cambridge Library Collection.
13. Tännsjö, T. (2020). *Why Derek Parfit had reasons to accept the repugnant conclusion*. *Utilitas*, 32(4), 387–397. <https://doi.org/10.1017/s0953820820000102>
14. Temkin, L. S. (2011). *Rethinking the good: Moral ideals and the nature of practical reasoning*. Oxford University Press.