Philosophia
https://doi.org/10.1007/s11406-022-00589-7

# Circumventing the Non-identity Problem

**Brian Carey**[1]

Received: 13 April 2022 / Revised: 28 September 2022 / Accepted: 18 October 2022


**Abstract**
This article provides an account of moral obligations that we have towards present generations, which require us to produce outcomes that are similar to those we would be required to produce if we had moral obligations to future generations. Discharging these duties enables us to secure the kinds of goods for future generations that we intuitively think we ought to provide in the absence of an answer to the non-identity problem. In this sense, the non- identity problem is avoided rather than solved. Nevertheless, a significant upshot of this account is that it provides a basis for practical action in the face of theoretical uncertainty.

**Keywords** The Non-identity Problem · Future Generations · Intergenerational Justice · Population Ethics

## 1 Introduction

Many of the decisions we make in the present will have profound effects upon future generations; determining whether the lives of future people go better or worse, as well as determining which people and how many people will exist. Many of us believe that this means there are moral obligations that we in the present bear towards those yet to be born. Specifying the structure, scope, and content of these obligations is obviously an important task. However, determining exactly what we owe to future generations has proved to be very difficult, thanks in large part to what has become known as the 'non-identity problem'.[1] The non-identity problem emerges when we combine the fact that our actions in the present determine which future people will exist with the intuitive view that what it means to harm a person is to make her life go worse than it otherwise would have. Thus, if we imagine a

[1] Early and influential discussions of this problem can be found in Kavka (1982), Parfit (1982, 1984), and Woodward (1986).

✉ Brian Carey
bpdcarey@gmail.com; brian.carey@durham.ac.uk

[1] School of Government & International Affairs, Durham University, Durham, UK

Published online: 23 November 2022

person born in the distant future, it is difficult to explain on what grounds she might complain that she has been wronged by the behavior of her ancestors if her very existence is predicated upon that same behavior.

In this article, I provide an account of particular moral obligations that we have towards certain members of *existing* generations, which require us to produce the same or very similar outcomes to those we *would* be obliged to produce *if* we had obligations to members of future generations who do not yet exist. One significant advantage of this account is that recognizing these obligations allows us to secure benefits for future generations that we intuitively think we ought to provide, without relying on claims about what we owe directly to members of future generations themselves.[2] This approach aims to avoid the non-identity problem, rather than solve it. However, so long as the non-identity problem remains a problem, the approach I will describe in this article is capable of providing us with sufficient reasons to act *as though* we have obligations to future generations, even in the absence of a consensus among philosophers about whether such obligations exist and how they ought to be understood.

In Section 2, I explain the non-identity problem in more detail and show why it is difficult to explain obligations towards members of future generations. In Section 3, I introduce the idea of 'moral overdetermination', and suggest that this concept can sometimes allow us to make practical progress in spite of deep philosophical disagreement about our moral obligations. In Section 4, I explain how the concept of moral overdetermination can help us to avoid the non-identity problem. I begin by identifying a particular set of obligations that we have to some members of presently-existing generations, and argue that these require us to pursue outcomes that are essentially the same as those we would be required to pursue if we did have obligations towards members of future generations. In Section 5, I discuss some of the limitations of this approach, given that it almost certainly does not capture our strongest reasons for securing a minimally decent future for those yet to be born. Section 6 concludes.

## 2 The Non-Identity Problem

Most of us have strong intuitions to the effect that it would be wrong for people living in the present to make policy decisions without any regard for the impact that those decisions might have on the lives of those yet to be born. Typical examples that philosophers use to illustrate such intuitions involve some seemingly selfish or reckless act in the present, which has dire consequences for the lives of future

[2] Axel Gosseries (2008) identifies an approach to this problem which he calls the "present rights of present persons" option, which is the type of solution I will argue for in this section. Howarth (1992) describes a "chain of obligation" between overlapping generations as part of a view which is similar to the solution described in this article (though there are some significant differences, such as my account of the precise content of our obligations to present generations and the distinctly non-ideal flavour of the solution as I shall describe it). See also, Steiner & Vallentyne (2009) on the issue of overlapping generations and conceptual theories of rights and Gardiner (2003) on why a form of the 'Pure Intergenerational Problem' persists even in cases involving overlapping generations.

generations. So, for example, perhaps a particular society decides to save money in the present by pursuing policies that they know will have catastrophic effects upon the environment, but these effects will only emerge long after the present generation have passed away.

The non-identity problem begins to emerge once we notice that many (if not all) of the people who will suffer in the future *are not the same people* as those who would enjoy the benefits of a clean environment in the future, had we acted differently in the present. This claim follows from any plausible conception of personal identity, such as a 'genetic identity' view of personhood for example, according to which an individual's existence depends (among other things) on the fact that they are the product of a particular sperm and egg.[3] In other words, had your parents chosen to conceive at a different time, or indeed with different people, you would never have existed. If we accept this (or a similar) view of personal identity, it is plausible to suppose that our behavior in the present – crucially, including our morally dubious behavior – affects *which* people will come into existence in the future.

Now suppose that we choose the selfish or reckless policy in the present, and this leads to a future where some particular person is born and suffers ill-health because of the polluted environment she has inherited. Intuitively, we are inclined to think she has some complaint against us, but it is difficult to pin down the exact nature of this complaint. Usually when we are wronged we can appeal to a *counterfactual account of harm*. Roughly put, such an account understands harm in terms of a comparison between how a person's life goes in reality, and how it would have gone, if not for the action(s) or inaction(s) of some other agent(s). For instance, if you steal some money from me, I can explain how you have harmed me by appealing to the counterfactual scenario in which you did not steal from me, and I was able to use my money to advance my own interests. Clearly the account as I have sketched it here is in need of some deeper philosophical scrutiny, but this basic version will do for our purposes since it seems to capture our ordinary intuitions about what it means to harm someone. The problem for our future person, whose life is going poorly as a result of her polluted environment is that *she simply would not exist* in a counterfactual scenario where her ancestors chose not to pollute. Thus she cannot say "if only you had not behaved in this way, my life would have been so much better", and must instead find some other grounds for complaint.

One option might be to find a way to understand the reckless or selfish behavior as being morally wrong in a way that does not depend on the lives of actual future people.

Such an approach runs contrary to another commonly-held view, however, which is that in order for an action to be wrong *there must be someone who is wronged*. This is the 'person- affecting principle'.[4] Those who accept such a principle hold that it simply wouldn't make sense to say that some act or omission could be morally wrong, without it thereby also being bad for at least one actual person.

[3] See Temkin (1987), Parfit (1984, p. 352) and Kripke (1980, pp. 111–116). For recent criticism of this conception, see Wrigley (2012).

[4] See Parfit (1984, p.363) and (2017).

Combining these different ideas – a genetic identity view of personal identity (or a view about personal identity with similar implications), a counterfactual account of harm, and the person-affecting principle – allows us to fully grasp the non-identity problem, and to see the tension between our commonly-held views about our duties toward other people in general, and our intuitions about what we owe to future people in particular. As I have sketched these issues above, there are a number of obvious points of attack for the philosopher attempting to offer a solution to this puzzle. Philosophers have argued for rejecting or modifying a genetic identity view (Wrigley, 2012), the counterfactual account of harm (Hanser, 2008; Shiffrin, 1999), and the person-affecting principle (Meacham, 2012; Ross, 2015). Despite these and similar attempts, no solution (yet) commands a consensus among philosophers as providing a plausible and comprehensive solution to the non-identity problem.[5]

Some theorists have instead attempted to avoid the non-identity problem rather than to solve it – this article is one such attempt, so it will be helpful before proceeding further to explain how the position I will develop here differs from other efforts to avoid the problem. The most significant difference between my approach and others that aim to avoid the non-identity problem is that mine is motivated not just by the problem itself but by the apparent lack of consensus as to whether or how we should aim to solve or avoid it. One implication of this approach is that there is a provisional aspect to my position that is not typically found in other attempts to avoid the non-identity problem: my view leaves open the possibility that a solution will be found or indeed that a solution has already been found (but not one that commands sufficient consensus among philosophers). A related implication is that the obligations I will identify below are entirely compatible not just with other attempts to avoid the non-identity problem but with other attempts to solve it. For this reason, the potential for consensus is greater than it is for other views that aim to avoid the non-identity problem – those views require us to give up on a solution to the problem, whereas mine does not.

If we want to produce action-guiding recommendations for people in the real world, we do not have the luxury of waiting for a consensus to emerge among philosophers when it comes to policy matters that have the potential to profoundly affect the lives of future generations. In the absence of a consensus, how ought we to proceed when the stakes are so high?

## 3 'Moral Overdetermination' and Moral Uncertainty

When we consider philosophical questions that have the potential to significantly impact people's lives, and when time is of the essence, a lack of a consensus among the relevant experts constitutes a distinct problem that requires a distinct solution. It is one thing to ask "what (if anything) do we owe to future generations?", but it is quite another to ask "what should we do in the absence of a consensus about what

[5] For an argument that we may have "partially" solved the problem, see Woollard (2012). For an argument that the non-identity problem has little practical relevance for climate policy, see Tank (2022).

(if anything) we owe to future generations?". My aim in this article is to offer an answer to the latter question, rather than the former. My suggestion is that one useful way to approach this type of question is to consider a phenomenon I shall call 'moral overdetermination'.

Here is a simple case of moral overdetermination. Suppose that Albert is a teenager who promises each of his parents separately that he will do his homework before dinner. If we think that promising either parent alone would have been sufficient to generate a duty for Albert to do his homework before dinner, we have a case of moral overdetermination. Albert possesses more than one morally weighty reason to perform a particular act, any of which would be sufficient on its own to provide him with a decisive reason to perform the act.

In the example just described, the strength of each reason is equal, but this need not be the case. Suppose that Albert is considering whether to steal from Betty. Upon reflection, Albert comes to believe that he has a good reason not to steal. This reason emerges from consideration of Betty's moral status – Albert concludes that Betty has a moral right to her property and that it would be deeply disrespectful of her moral status if Albert were to ignore this (if this strikes you as not quite right, substitute your preferred account of the immorality of theft instead). Let us say that this reason is sufficiently powerful such that it is capable of justifying Albert's decision not to steal from Betty. Upon further reflection, however, Albert also recalls a promise that he made to his friend Charlie to the effect that he would not steal from Betty. Were he to steal from Betty, Albert would be breaking his promise, and Albert believes it is usually wrong to break your promises. Let us say that this reason is also sufficiently powerful such that it would motivate Albert to refrain from stealing even if he had not also recognized that such an act would also violate his duties toward Betty. What makes each of these examples a case of 'moral overdetermination' is that Albert has more than one moral duty, each of which can be discharged via the same (in)action, and each of which would be sufficient on its own to justify the (in)action in question (in this case, to refrain from stealing from Betty).

The key point to note here is that moral overdetermination refers to the reasons that an agent has to act regardless of whether or to what extent the agent recognizes those reasons. All that matters is that the reasons are individually sufficient to justify action, and that they are jointly realizable. Among all such reasons, there may be many of which we are not aware or which we do not fully understand. In cases where moral overdetermination applies, there is no reason to assume that we will necessarily have a better awareness or understanding of stronger reasons rather than weaker ones.

This may seem counterintuitive. After all, in the second case described above it seems strange to imagine a person who would be aware that it's wrong to break promises but who wouldn't be aware that it's wrong to steal. This point can be answered by distinguishing between having the belief that something is wrong and being able to fully explain why it is so. Consider any debate in moral philosophy where philosophers agree with a particular conclusion but dispute the means of getting there. Alternatively, we might simply note that there is no necessary conceptual connection between the complexity of a moral question, and the strength of the

moral reasons it may produce. Trivial questions may be difficult to answer, and serious questions may be easy to answer.

Unfortunately, the subject of this article is a serious question without an easy answer. The non-identity problem has yet to be solved – or at least, no solution yet commands a consensus among the majority of philosophers who have sought to answer it. However, might this be a case where moral overdetermination applies? In the absence of a consensus-commanding explanation as to why we could wrong future people by pursuing certain policies in the present, might we be able to access other reasons, lesser in strength but nevertheless sufficient to recommend a course of action similar to that which would be recommended by any plausible solution to the non- identity problem?

It should be clear at this point that an answer in the affirmative will necessarily be unsatisfying in an important sense. It will be unsatisfying for the same reason that it is unsatisfying to imagine Albert choosing not to steal from Betty only because he has made a promise to Charlie. We don't just want Albert to be motivated by a sufficiently good reason – we want Albert to be motivated by the *strongest* reason. Nevertheless, if Albert is presently unable to figure out what that is, we should at least think it better that he be motivated by the weaker one than by none at all. Similarly, if we are presently unable to explain the structure and content of our moral obligations toward future generations, we can at least hope to find some weaker reasons that are nevertheless powerful enough to motivate us to act as though such obligations did exist. In the next section, I offer one possible source of such reasons.

Before proceeding, there is one preliminary objection that must be dealt with. In my characterization of moral overdetermination in this section, I have said that moral overdetermination requires that each of the relevant reasons be jointly realizable. This feature is important because I will argue in what follows that we should act on the basis of reasons that are likely to be *weaker* than other reasons which we will (hopefully) be able to access (or reach a consensus upon) in the future. If these stronger reasons were to contradict the weaker reasons, then it would be an open question as to whether we should act on the reasons available to us, or wait to see whether we could identify and build a consensus around the stronger reasons at some point in the future. However, if we have good reason to think that the weaker (and accessible) reasons are jointly realizable with the stronger reasons, then we can safely act on the basis of the weaker reasons without fear of undermining the stronger.

In the case described above, Albert's decision not to steal from Betty can satisfy both his duty to Betty and his (weaker) duty to Charlie. However, it might be objected that Albert can safely act on the basis of the weaker duty only if he knows that it is jointly realizable with any stronger duties he may have. Thus, so the objection might go, Albert can only safely act on the basis of the weaker reason if he can also perceive the stronger reason. Applying this conclusion to the case of the non-identity problem, we might think that the concept of moral overdetermination cannot help us act on the basis of incomplete knowledge of moral facts, since knowing that we are in a situation where moral overdetermination applies presupposes that we already know all of the relevant moral facts.

To answer this objection, we should look more closely at what Albert needs to believe in order to be motivated by the weaker reason. Obviously, Albert needs to believe that the weaker reason is sufficient to demand action unless contradicted by a stronger reason. The crux of the issue will be Albert's belief that the weaker reason is not in fact (or is highly unlikely to be) contradicted by a stronger reason, even though Albert (let us stipulate) cannot explain to us precisely what that stronger reason is. Here it will be helpful to deploy the distinction made earlier between having the intuition that we have a duty to perform a particular action and the ability to explain that intuition in rigorous philosophical terms. For someone in Albert's epistemic position, the suspicion that moral overdetermination may apply to his case can arise only if he has the intuition that there is some stronger reason that should motivate his actions, though he cannot fully account for it. From Albert's perspective, he knows what his strongest reasons say he ought to do, but cannot fully explain why he ought to do it.

Under the circumstances just described, it seems clear that Albert is entitled, if not required, to act on the basis of the weaker reason. This is because Albert *is* in an epistemic position to act as though his actions won't be overridden by a stronger reason, since he knows *enough* about the stronger reason to know that its recommendations do not conflict with those of the weaker. Thus we should conclude that some level of awareness of one's reasons is necessary in order to act as though one is involved in a case of moral overdetermination and that this level of awareness requires only that we know what our strongest reasons require, but not why.

Applying the above considerations to the non-identity problem, the arguments to follow will work best if the reader shares my intuitions that there must be strong reasons for us not to adopt policies in the present that will have disastrous effects upon the lives of future generations and that these reasons connect us to future people such that we can wrong those people themselves by our actions in the present. However, if the reader has no such intuitions (either because you believe that future generations have no moral status, or because you believe we cannot wrong future people with our actions), then the reasons to be discussed below may in fact turn out to be stronger than I imagine them to be.

## 4 Justice for Future Generations via Justice for Present Generations

Suppose that we have a set consisting of all the possible children that some particular parent might raise[6] over the course of their life. Within this set, suppose we have a subset of possible children who will lead lives worth living, and within that subset a further subset of possible children who will lead lives that are not just worth living, but who will have a good chance of living a fully flourishing life.

This proposed terminology requires some clarification before we can proceed.[7] Let us say that a person's life is "flourishing" to the extent that they successfully

[6] In what follows, I refer to parents "raising" children, rather than "having" children so that my account is inclusive of parent/child relationships beyond those of a biological parent and child.

[7] I am grateful to an anonymous reviewer for pressing me to clarify these terms.

realize the projects that are most valuable for them (I shall refer to these simply as their "projects" in what follows). The sense of "value" in question here may be understood objectively or subjectively[8] but nothing in the arguments to follow will depend on which approach we take on that question. This conception of flourishing does not assume that one's projects are aggregable or comparable, just that a person's life will tend to go better if as many such projects as possible are realized[9] to the extent that these projects are jointly-realizable and pursuable in ways that are morally permissible. It may be unlikely that many (if any) people in the real world lead a fully flourishing life in the sense that all of their projects are realized, but it seems clear that the ideal of a fully flourishing life is something that many (if not most) of us pursue, and that (in)actions by third parties that may undermine our ability to pursue a fully flourishing life stand in need of justification.

Let us say that one has a "good chance" of living a fully flourishing life to the extent that, if one were to try and not give up, one would tend to succeed in realizing as many of one's projects as possible.[10] This is compatible with the fact that realizing some projects may make it difficult or impossible to realize others, and there may be projects we are morally prohibited from realizing, depending on the circumstances.

Let us call those possible children who would have a good chance of leading a fully flourishing life 'fully flourishing children'. It seems plausible to say that a parent has a strong interest in being able to choose to raise a fully flourishing child (specifically, in being able to choose to pursue this ideal for their child, even if it is never fully realized). This is not to say that parents *ought* to choose to raise a fully flourishing child, nor that their interest in having the choice can never be overridden by some more significant consideration. Nevertheless, having the capacity to be able to choose to raise a fully flourishing child represents a very significant interest for many, perhaps most, perhaps all, potential parents.

To support this claim, consider the case of a potential parent who is deciding whether to raise a child now, at a point in time when they have few resources available to offer the child, or to wait until some point in the future where they believe they will be in a better position to provide a better life for a different child.[11] Suppose that if they raise a child now, they know it will have a life worth living, but won't have a good chance of being a fully flourishing child, but if they wait they can raise a fully flourishing child. If some third party were to intervene to remove the option to wait, they would have damaged the parent's interests in an important way,

[8] A subjective understanding in this context is one in which "facts about a person's welfare depend on facts about her actual or hypothetical mental states" (Hawkins, 2010, p.62), while an objective understanding holds the relevant projects to be "intrinsically valuable independently of anyone's judgments, attitudes, desires, or dispositions" (Hewitt, 2010, p.344) concerning those projects.

[9] Different kinds of projects will have different kinds of success conditions, for example some projects may allow for partial or different degrees of success, and some projects may require specific conditions to be sustained over time while others may succeed or fail at a particular moment in a person's life.

[10] I borrow this wording from David Estlund's counterfactual conditional account of political feasibility. (Estlund, 2011, p.212).

[11] This is a version of a case described by Parfit (1984, p. 358).

even if they would have chosen to raise the child now rather than to wait.[12] It is crucial to emphasize this point that the arguments to follow do not depend on the claim (one I take to be implausible) that raising a fully flourishing child (or indeed, raising a child at all) ought to be among one's projects. Rather, the claim is that each of us has an interest in being able to choose which projects to pursue, and that this interest is undermined even when options are removed for us that we would not have chosen ourselves.

Thus, the interests in question here are not to be confused with the interest each of us has in ensuring that our loved ones live worthwhile lives (no child may exist at the point before the parent makes their decision), but rather an interest in being able to create and sustain worthwhile relationships in the future which may not exist at present, and the interest to choose whether to create and sustain such relationships. These interests are grounded in both the potential benefits that such relationships may bring, as well as the benefit of being able to autonomously choose what sort of relationships to create. Significantly, the project of creating a fully flourishing child is one that tends to be ubiquitous, in that it tends to be among many people's projects and non-compensatable, in that denying a person the opportunity to pursue such a project cannot simply be compensated for by allowing them to succeed in other projects. We may think that projects with these features are especially important to protect, since their ubiquity makes it more likely that people will tend to want to pursue them, and their non-compensatability makes it harder to rectify the harm of denying the opportunity to those who want it.

Nevertheless, one might object to the suggestion that the potential parent's interests are necessarily harmed if they are unable to choose whether to raise a child now or to wait. If they would have chosen to wait, so the objection might go, then the potential parent has not really had their interests undermined if the other option is removed. Indeed, if we endorse a counterfactual account of harm, how can it be said that they are harmed by removing their ability to choose something they would not have chosen otherwise?

At least two replies can be offered here. The first concerns the practical difficulties inherent in removing choice from a person on the basis that we are entitled to assume which decision they would have made had we not intervened. While certain goods may be presumptive goods, in that we are entitled to presume that all reasonable and rational people would desire them, it seems wrong to presume that reasonable and rational people will necessarily choose to wait, if given the choice faced by the potential parent in our example. We are not morally required to maximize the potential benefits that our future children will enjoy, only to take reasonable steps to ensure that we do not create children whose lives are not worth living. Provided that that threshold has been met, reasonable and rational people will disagree about whether and when to have children. To presume to make that decision on behalf of another is to impose our own will upon theirs, and should strike us as unjustifiably paternalistic on that basis. On the contrary, it seems far more plausible to say that

[12] Note that we can agree with this while remaining entirely agnostic as to whether it is permissible for the parent to have the child now, or whether they are required to wait.

reasonable and rational people will agree that it is right to protect a person's ability to make their own decisions concerning whether and when to raise children (including the ability to revise such decisions, where applicable).

This leads us to a second reply to the objection in question, which concerns the value of *choosing for oneself*. If we think that there is something valuable about the exercise of autonomous choice, then a counterfactual account of harm can easily explain how a person can be harmed by removing an option that they otherwise would not have chosen. Specifically, the harm that is caused in such cases is explained by the fact that the person in question was deprived of the opportunity to choose, and choosing is itself a morally valuable activity insofar as it constitutes the exercise of a person's autonomy. To emphasize the point, we can consider cases where a person is prevented from choosing something that is bad for them: suppose you prevent me from having a second helping of dessert because you (reasonably and correctly) believe it will be bad for my health. In such cases we might well think that I have been harmed in one sense (by being deprived of the ability to choose) even though I have been benefitted in another (by being deprived of another slice of some delicious but calorific chocolate cake). Thus the parent in our example may object that removing their ability to choose to raise a child now has harmed them, even though they would in fact have chosen to wait, if they had had the ability to choose. As above, in denying the potential parent the ability to exercise their autonomy, we impose our own will upon theirs, and act paternalistically toward them.

Bearing in mind that nothing in what I have argued or will argue suggests that anyone is under a moral obligation to choose to raise a fully flourishing child, let's call the capacity *to choose* to have a fully flourishing child Capacity X. If we alter a potential parent's range of options, such that we undermine or deprive them of Capacity X, we harm them to the extent that we make it less likely or impossible that *they* will be able to have a good chance of living a fully flourishing life. If we remove the parent's option to wait, for instance, this constitutes a harm to them, even if they would not have chosen to wait, and even if they would not have chosen to raise a child, provided that having that option open to them is valuable in itself.

Given this, it looks as though we have (at least) a prima facie reason not to hinder a person's capacity to choose to raise fully flourishing children. Furthermore, we can say that a person who lacks Capacity X, will themselves be unable to lead a flourishing life of their own, since at least one of their projects has been rendered unrealizable (i.e. the project of being able to choose whether to exercise Capacity X). In other words, one's chances of being a fully flourishing child are undermined if one lacks the capacity to raise fully flourishing children, if one chooses. Crucially, this is true regardless of whether one actually values exercising Capacity X – someone who does wish to raise a fully flourishing child and is prevented from doing so will have two of their projects undermined (their capacity to exercise autonomous choice plus the ability to pursue the option they would have chosen) and so their lives may go even worse than those who are prevented from choosing but who would have chosen not to raise a fully flourishing child. In either case, so long as a person's ability to choose has been undermined, and so long as we take the ability to exercise autonomous choice of this sort to be valuable for a person, their interests are undermined if Capacity X is undermined.

My focus here is on "fully" flourishing lives in order to avoid the stronger claim that Capacity X is necessary in order for one to enjoy very high levels of wellbeing – there is an alternative sense of "flourishing" that is binary rather than scalar, according to which we might say that a person lives a flourishing life if at least some (or enough) of their projects are realized. In this sense, of course we can imagine someone living a flourishing life who lacks Capacity X. Perhaps most obviously, there are elderly people who are unable to procreate and legally prohibited from adopting due to their age. We wouldn't want to say that such people cannot lead flourishing lives on that basis (nor indeed would we want to say that their lives cannot be "fully" flourishing either). We can reply to such an objection by insisting that what counts as a fully flourishing life needs to be understood holistically, such that what matters is whether the elderly people in question possessed Capacity X for some reasonable length of time at some point in their lives, and if not, whether the reason why they lack Capacity X is due to the blameworthy actions of others (in which case is perhaps more plausible to suggest that this would have undermined their ability to live a fully flourishing life).

None of the above entails that one cannot live a fully flourishing life without choosing to raise fully flourishing children. What matters here is the interests people have in having significant choices. The most controversial claim that *is* entailed by this view (which I hope is not particularly controversial) is that one's ability to lead a fully flourishing life is undermined if one is deprived (via the blameworthy actions of others) of the ability to choose whether to (try to) raise a fully flourishing child. To summarize, here are four claims that are central to my argument thus far:

(1) Being a fully flourishing child requires having a good chance of living a fully flourishing life.
(2) People generally have a strong interest in being able to choose whether to raise fully flourishing children [ to be able to do so is to possess Capacity X].
(3) People who lack Capacity X do not live fully flourishing lives, if possession of Capacity X is among their projects.

In addition to the above, I shall now argue for the following claims:

(4) To deny someone Capacity X is to undermine their ability to be a fully flourishing child.
(5) We have a strong reason to ensure that people possess Capacity X.

Given the first three claims, suppose that we have 26 overlapping generations of people, from Generations A, B, and C, which exist in the present, to Generations X, Y and Z, which will exist at some point in the future. Suppose that the adult members of existing generations adopt a policy which will see a rise in the standard of living for each subsequent generation until the point that members of Generation Y are born, at which point members of Generation Y will only be able to raise children

who lead lives worth living, but do not have a good chance of leading fully flourishing lives (because, if they were to try and not give up, they would not tend to succeed in realizing at least some of their projects).

It follows that members of Generation Y lack Capacity X, and (following (4) above) their ability to be fully flourishing has been undermined. If (many, most, or all of) the members of Generation Y lack Capacity X and if the capacity to raise fully flourishing children is important for one to lead a fully flourishing life, then we can assume that (many, most, or all of) Generation X also lack Capacity X. In other words, the fact that members of Generation Y cannot raise fully flourishing children means that they themselves are not fully flourishing children, which means that the preceding generation were not fully flourishing children either. We can tell the same story for the parents of Generation X, and so on, producing a regression which takes us all the way back to the original actions of the adult members of existing generations (say, Generations A and B), which deprive their existing children (Generation C) of Capacity X, thereby violating duties that the adult members of existing generations have to their children, and to their fellow adults.

Crucially, this conclusion is entirely compatible with the counterfactual account of harm and the person-affecting principle. In choosing the policy in question, members of Generations A and B deprive their children of Capacity X, rendering them unable to pursue fully flourishing lives as a result. According to the argument in the preceding paragraph, this also has of the same effect on Generations A and B themselves, but we can focus on their duties towards children who are members of Generation C for the sake of simplicity.[13] Recall that the counterfactual account of harm requires us to show that a person is worse off than they otherwise would have been due to the (in)actions of another, and that this presents a problem for obligations to future generations because their very existence may be contingent upon 'harmful' acts we take in the present. It is difficult to know how to square this account with the status of future generations, if we assume that they lead lives that are sub-optimal yet preferable to non-existence. However, in the case just described, the existence of the members of Generation C is *not* contingent on whether Generations A and B choose the policy in question, thus the counterfactual account of harm can be applied.

Recall also that the person-affecting principle says that in order for an (in)action to count as wrong it must be wrong *for someone*, and that this presents a challenge for accounts of our obligations to future generations whose existence depends on the (in)actions in question. In the case just described, however, the person-affecting principle is easily satisfied: members of Generation C are alive here and now. Choosing the policy in question is wrong for *those people*, regardless of whether it is also wrong for members of future generations. Yet the practical requirement of vindicating our duties to younger (existing) generations is that we adults must ensure (to the best of our abilities, within reason) that our actions in the present do not impose suffering upon those yet to be born. This is the only way to preserve Capacity X for ourselves and members of the younger, overlapping, generation.

[13] Given that most members of Generation C are children and thus not full moral agents, this further simplifies the picture, given that we cannot (yet) attribute any culpability to members of Generation C for the situation they find themselves in.

One might object that the case as originally described involves pursuing a policy that improves the lives of existing generations (and several subsequent generations), before the harmful consequences are felt by distant future generations. What if the benefits to the members of Generation C are so significant that they would outweigh the loss of Capacity X? In such cases we may still have the intuition that the policy ought not to be pursued (because of its harmful effects upon distant generations), but we could no longer explain why the policy should be avoided in terms of obligations to presently-existing generations.

In response, we should concede that trade-offs might be justified in cases where the benefits of the policy in question are necessary to secure some other important good that significantly affects a person's ability to live a flourishing life. In such cases, however, our intuitions about our obligations to future generations are likely to change accordingly. There is an obvious difference between a policy designed, for example, to provide luxurious living standards to presently-living people which are well above what is necessary to lead a decent life, and a policy designed to secure access to basic goods like food, healthcare, or education. Our intuitions that we wrong future generations are strongest when the policies in question serve to push existing generations far above any plausible minimum threshold of well-being (and/or to push future generations below it). It is important to concede that present generations are not required to sacrifice absolutely everything for the needs of future generations, and to acknowledge that Capacity X is just one of many important factors that contributes to a fully flourishing life.

Another possible objection to this approach concerns the epistemic constraints associated with making assumptions about what would count as a fully flourishing life for members of future generations. According to this objection, what should count as a fully flourishing life is likely to change over time, and we are not in a position to predict whether policies we choose today will contribute to, detract from, or make no meaningful difference to whether members of future generations are able to lead fully flourishing lives. The correct response to this objection is to concede the point that the content of a fully flourishing life may change over time, but to deny that this gets us off the hook when considering which policies to pursue in the present. One reason for this is the fact that certain goods do seem to persist through time, such that a world where they were not part of a fully flourishing life would look very different to our own. These include, for example, health, liberty, education, and the capacity to form and sustain meaningful social and intimate relationships with one another (as well as the good of being able to autonomously choose how we value and pursue projects related to these). It is reasonable for us to presume that such goods will continue to be ingredients of a fully flourishing life well into the future.

Of course, there remains the possibility that over time what counts as a fully flourishing life will change in extreme and/or unforeseeable ways. Two points are worth making here. First, while those of us in the present generation wield extraordinary power when it comes to determining which future people will be born and how their lives will go, this power is not limitless, nor concentrated entirely within our own generation. Our obligations to distant generations will be shared with generations in between, and (with certain important exceptions involving the use of finite resources in the present) our ability to affect the lives of future generations will diminish over

time, through no fault of our own. Second, our ability to foresee what future generations will need and how to secure it will also diminish over time, through no fault of our own. Accepting these claims yields a picture of our obligations to future generations which suggests they become weaker over time, as our abilities to affect the future diminish. However, none of these considerations can ground a justification for inaction in the present, given what we do know and what we can change about the likely lives of future generations who will exist in the not-too-distant future.

## 5 Circumventing the Non-Identity Problem

The approach set out in the previous sections is one which can provide us with reasons for action, but almost certainly fails to (and is not an attempt to) identify the strongest reasons that we have to avoid harming members of future generations. In this section my aim is to explicitly acknowledge and explore the limitations inherent in this approach, but in so doing to demonstrate the value of having access to sufficient reasons to act, even while we lack access to our best reasons.

One obvious limitation of my approach concerns the account of human flourishing on which my arguments depend. In the previous sections, I have tried to make as few claims as possible about what constitutes a flourishing life. One risk here is that an overly narrow account of human flourishing may generate a new problem of precisely the sort that this article is intended to avoid, i.e. that the arguments for circumventing the non-identity problem may come to depend upon a controversial account of human flourishing about which there is no general consensus. On the other hand, if we adopt an understanding of human flourishing that is especially broad in scope, then it becomes more difficult to see why Capacity X (or perhaps almost any capacity) should necessarily or presumptively count as part of a fully flourishing life.

My preferred reply to this challenge is to insist that there really is a general consensus that Capacity X should be thought of as a presumptive good, provided that Capacity X is properly specified. Recall that Capacity X tracks an ability to choose whether to pursue the project of raising a fully flourishing child. It is not controversial to say that individuals have a strong interest in being able to choose the kinds of relationships they form, and it is plausible to think that this interest is strongest when it comes to familial relationships. It is very bad for a person if their ability to make friends or acquaintances is constrained, but it would tend to be even worse for them if they were deprived of the ability to form, sustain, and shape relationships with romantic partners, parents, or children. In order to identify a consensus around the general importance of Capacity X, we need only to recognize that having the ability to make these decisions matters a lot for lots of people. All of the above is entirely consistent with the view that a person may lead a fully flourishing life without forming many or even any such relationships, and with the view that worthwhile relationships come in a wide range of forms.

This last point may seem contradictory: if possession of Capacity X tends to be important for a person to lead a fully flourishing life, how can I claim both that parents need not choose to raise fully flourishing children, and that parents have a duty not to deprive their children of Capacity X? Surely something has to give here: either we bite the bullet and insist that

parents are morally required to (try their best to) raise fully flourishing children or we concede that it is fine for parents to deprive their children (and future generations) of Capacity X.

In reply, it is worth emphasizing first that the arguments in the preceding sections do not entail that possession of Capacity X *is the same thing* as being a fully flourishing child. To be a fully flourishing child is to have a good chance of leading a life where none of one's most significant projects are thwarted – on any plausible account of flourishing lives, this will include far more than just the ability to exercise Capacity X. So I could say that my account does not entail that parents must choose to raise fully flourishing children, but it does perhaps entail that parents should try to ensure that their children possess Capacity X. This would be too strong a claim, however – the obligations discussed throughout this article are *pro tanto* obligations. They give rise to reasons that are defeasible by other considerations, provided that those other considerations are strong enough. A parent should not be required to raise a fully flourishing child if doing so would deprive her of some even more important good (such as important aspects of her own autonomy, for instance).

A simpler reply is available, which is to distinguish between a parent's duties when considering whether and how to raise a child and a parent's duties to children after the relationship has been established. Setting aside cases where a parent may be obliged not to create a child whose life is not worth living, we can insist that parents who possess Capacity X and choose to raise children are obliged to do what they can (within reason) to secure Capacity X for their children once those relationships have been established. Such an approach allows us to maintain obligations to present (and future) generations without infringing upon parents' autonomy.

The preceding paragraphs are intended to show that the general approach argued for throughout this article does not rely upon controversial or counterintuitive claims about the good in general, or about parents' obligations toward their children in particular. This is especially important in this context, given that I aim to identify a basis for action which should be capable of commanding a consensus among most reasonable people (in contrast to theories that aim to specify our obligations to future generations). If the obligations I have argued for here are no better at generating a consensus, then we are back where we began. On the other hand, if I have set out a convincing case for our obligations to secure Capacity X for presently existing generations, then we have good reasons (albeit not the best reasons) to act in ways that will benefit those who have yet to be born. Such reasons can provide sufficient justification for action here and now, until a broader consensus can be reached as to the best explanation for our obligations to future generations.

## 6 Conclusion

In this article, my central aim has been to describe and defend moral duties that we have toward existing members of younger generations, whose fulfilment is likely to secure what we intuitively think we owe to future generations. I suggested that this is a case of 'moral overdetermination', such that we are able to fully specify some of the reasons to bring about a particular outcome (ensuring that present and future generations possess Capacity X), despite not having access to (or not having a consensus about) the strongest reasons for action. The upshot of this approach, I have argued, is that it can explain why we ought to provide future generations with the opportunity

to lead fully flourishing lives, while avoiding contentious and difficult philosophical problems such as the non-identity problem and person-affecting principle, as well as the fact that there is no general consensus among the experts as to the best explanation for our intuitions regarding our moral duties toward future generations.

## Declarations

**Conflicts of Interest/Competing Interests** The author declares none.



## References

Estlund, D. (2011). Human nature and the limits (If Any) of political philosophy. *Philosophy & Public Affairs, 39*(3), 207–237.
Gardiner, S. M. (2003). The pure intergenerational problem. *The Monist, 83*(3), 481–500.
Gosseries, A. (2008). On future generations' future rights. *The Journal of Political Philosophy, 16*(4), 446–474.
Hanser, M. (2008). The metaphysics of harm. *Philosophy and Phenomenological Research, 77*(2), 421–450.
Hawkins, J. S. (2010). The subjective Intuition. *Philosophical Studies, 148*(1), 61–68.
Hewitt, S. (2010). What do our intuitions about the experience machine really tell us about hedonism? *Philosoiphical Studies, 151*(3), 331–349.
Howarth, R. B. (1992). Intergenerational justice and the chain of obligation. *Environmental Values, 1*(2), 133–140.
Kavka, G. (1982). The paradox of future individuals. *Philosophy & Public Affairs, 11*, 93–112.
Kripke, S. (1980). *Naming and necessity*. Harvard University Press.
Meacham, C. J. G. (2012). Person-affecting views and saturating counterpart relations. *Philosophical Studies, 158*(2), 257–287.
Parfit, D. (1982). Future generations: Further problems. *Philosophy & Public Affairs, 11*(2), 113–172.
Parfit, D. (1984). *Reasons and Persons*. Clarendon Press.
Parfit, D. (2017). Future people, the non-identity problem, and person-affecting principles. *Philosophy & Public Affairs, 45*(2), 118–157.
Ross, J. (2015). Rethinking the person-affecting principle. *Journal of Moral Philosophy, 12*(4), 428–461.
Shiffrin, S. (1999). Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory, 5*, 117–148.
Steiner, H. & Vallentyne, P. (2009). Libertarian theories of intergenerational justice. In A. Gosseries & L. Meyer (Eds.), *Justice between generations* (pp. 50–76). Oxford University Press.
Tank, L. (2022). Climate change and non-identity. *Utilitas, 34*(1), 84–96.
Temkin, L. S. (1987). Intransitivity and the mere addition paradox. *Philosophy and Public Affairs, 16*(2), 138–187.
Woodward, J. (1986). The non-identity problem. *Ethics, 96*(4), 804–831.
Woollard, F. (2012). Have we solved the non-identity problem? *Ethical Theory and Moral Practice, 15*(5), 677–690.
Wrigley, A. (2012). Harm to future persons: Non-identity problems and counterpart solutions. *Ethical Theory and Moral Practice, 15*, 175–190.