

Working paper 2026:10

Improving Lives and Avoiding Harm

by Tim Campbell

Improving Lives and Avoiding Harm

Abstract

The Procreation Asymmetry holds that creating a miserable person is morally impermissible, while not creating a happy person is morally permissible. Harm-avoidance theories explain this asymmetry by claiming that, in certain contexts, only actions that result in harm are impermissible. Two challenges for these theories are: (i) The Problem of Improvable Life Avoidance—requiring agents to avoid creating improvable good lives—and (ii) requiring dominated options, those that are worse for some and better for none. Philosophers have developed sophisticated harm-avoidance theories that address both challenges. However, we prove that *any* harm-avoidance theory entails either: (1) a worse form of The Problem of Improvable Life Avoidance, (2) an absurd form of sadism, (3) the permissibility of allowing unlimited harm, or (4) moral dilemmas. This result challenges both harm-avoidance theories and the Procreation Asymmetry, since it is doubtful that any alternative theory could support this asymmetry while avoiding (1)–(4).

Keywords: Population ethics, Procreation Asymmetry, Harm-avoidance theories, Problem of Improvable Life Avoidance, Impossibility theorems

1 Introduction

Suppose that if a person were to exist, she would be utterly miserable. She would have negative well-being, meaning her life would be bad in absolute terms. Could this make it wrong to create her? Most say ‘yes’. What if her life would instead be *good*? Could this, by itself,

make it wrong *not* to create her? Many who say ‘yes’ to the first question say ‘no’ to the second. According to them, creating a person with negative well-being is morally impermissible but refraining from creating a person with positive well-being is morally permissible, *all else being equal*.¹ This is *The Procreation Asymmetry*, or ‘the Asymmetry’ for short.²

What explains the Asymmetry? A common answer is that creating a person with negative well-being results in *harm* to that person, but failing to create a person with positive well-being does not result in harm to her, since there is no *her*.³ Call this *the harm-avoidance explanation* of the Asymmetry.

More generally, some defend what we call *harm-avoidance theories*. These theories say that in certain choice situations, which we outline in §2, an act is impermissible only if it results in harm.⁴ This explains why one is not required to create a person with a good life, all else being equal, since one cannot harm a non-existent person. But harm-avoidance theories can also provide a harm-based explanation of the impermissibility of creating people with bad lives. They can fully explain the Asymmetry.

However, it is unclear exactly how harm-avoidance theories should be spelled out. The simplest harm-avoidance theories, which require agents to minimize total harm, face *The Problem of Improvable Life Avoidance*, and morally require dominated options. The search for a harm-avoidance theory that deals adequately with these problems is underway.

We argue that this search is ill-starred. Any harm-avoidance theory either faces a more troubling variant of The Problem of Improvable Life Avoidance, permits agents to create people with miserable lives rather than people with good lives, permits acts that result in arbitrarily more harm for each of arbitrarily many more people than the alternatives, or accepts moral dilemmas, choice-contexts in which every option is impermissible.

We begin, in §2, by explaining what defines harm-avoidance theories, identifying a central commitment of such theories that we call *Harmless Permission* (cf. §2.4). In §3, we present The Problem of Improvable Life Avoidance and the problem of requiring dominated options. In §§4–5, we scrutinize two state-of-the-art harm-avoidance theories put forward by

Michael McDermott (2019) and Joe Horton (2021). While these theories evade The Problem of Improvable Life Avoidance, and do not require dominated options, they encounter at least one of the four other, no less severe, problems listed above. Corresponding to each of these problems, we identify four plausible principles. In §6, we prove that no harm-avoidance theory can accommodate all four. We conclude in §7, explaining how our challenge to harm-avoidance theories casts doubt on the Asymmetry.

2 Harm-avoidance theories

According to harm-avoidance theories, in certain choice situations, an option is impermissible only if it results in harm. In this section, we say what these choice situations are, what ‘harm’ means, and what counts as a harm-avoidance theory.

There are four points of clarification regarding harm-avoidance theories: domain of application (§2.1); relevant notion of ‘harm’ (§2.2); the conceptual framings these theories can adopt (§2.3); and how they address *non-identity cases*, in which the agent can choose between creating different (non-identical) people (§2.4).

2.1 Domain restriction

Some philosophers, such as McDermott (2019), Horton (2021), Thomas (2022), and Podgorski (2023) have proposed general theories that we shall classify as harm-avoidance theories. However, these theories either explicitly include (in the case of Thomas 2022), or are charitably interpreted as including (in the case of McDermott 2019, Horton 2021, and Podgorski 2023) a domain restriction. Specifically, we assume all harm-avoidance theories are restricted to the domain of *normative population ethics*—the part of moral philosophy concerned with acts that affect the (i) number, (ii) identities, and (iii) well-being levels of people, and with how considerations (i)–(iii) affect permissibility.⁵ For our purposes, considerations other than (i)–(iii), such as personal virtue, agent-relative prerogatives, special obligations,

and whether agents lie, cheat, and steal are sidelined. In the cases we shall consider, such considerations are not in play.

2.2 Harm

Another clarification concerns ‘harm’. Harm-avoidance theories typically recognize only two types of harm as morally significant: comparative and existential harm.

Suppose you have options A and B. If a person would exist under each option, and is worse off under A, then A results in comparative harm to her. If she has negative well-being under A but does not exist under B, then A results in existential harm to her.⁶

We deliberately use the phrase ‘results in harm’ rather than ‘harms’. If I allow a child to drown when I could easily rescue her at no cost, then she suffers a comparative harm—the harm of death. However, some would deny that I thereby *harm* her. They would say I merely *allow* harm to befall her. What matters on the theories we consider is the harm incurred by individuals under different options, regardless of whether these options may be said to *harm* them.

The harm resulting from an option is defined as comparative or existential relative to an alternative.⁷ It is therefore possible that a person suffers comparative harm under some option A relative to some alternative B, but also existential harm relative to another alternative, C. For instance, suppose I face the decision in Table 1. The well-being levels of people affected by my choice are represented numerically, where positive and negative numbers represent, respectively, positive and negative well-being, 0 represents neutral well-being, and ‘ Ω ’ represents non-existence.⁸

	Anne
Option 1	Ω
Option 2	-10
Option 3	10

Table 1: Improvable Life Case

Anne has an *improvable life* under Option 2, meaning she suffers comparative harm under that option.⁹ This involves a comparison of Anne’s well-being under Options 2 and 3. But Anne also suffers existential under Option 2, relative to Option 1.

This raises the question: *how much* harm does Anne suffer under Option 2? One possibility is that the harm she suffers has different magnitudes relative to different alternatives, but no magnitude full stop. Another is that the magnitude of Anne’s harm full stop is the difference between her well-being under Option 2 and her well-being under Option 3, which is the option that is best for her. When multiple individuals suffer harm under an option, one can either treat the total harm suffered (the sum of individual harms) as varying relative to alternatives, or else provide a criterion for total harm full stop. We shall return to this issue in §3.

2.3 Conceptual framings

We adopt the language of ‘harm’ to describe the class of theories that concerns us. But there is no unique conceptual framing that all these theories have in common. For instance, while Thomas (2022) states his theories in terms of harm, Horton (2021) states his in terms of *objections* to an agent’s behavior, and McDermott (2019) in terms of both harm and objections.¹⁰ Yet *the basis* of an individual’s objection, on these theories, is the existential or comparative harm she incurs as a result of an agent’s behavior. More generally, different conceptual framings are adopted by different philosophers who either explicitly endorse, or appear to

be committed to, the harm-avoidance explanation of the Asymmetry.¹¹

2.4 Non-identity Cases

The final point of clarification concerns *the Non-identity Problem*. In non-identity cases, an agent has at least two options, A and B, where A would create someone who would not exist under B, and B would create someone who would not exist under A.¹² These cases involve a choice between creating different *choice-dependent* people, i.e., people whose existence depends on the agent’s choice.¹³ See Table 2.

	Bob	Carl
Option 1	1	Ω
Option 2	Ω	10

Table 2: Non-identity Case

In contrast to non-identity cases, *addition cases* are those in which an agent can choose whether to create some choice-dependent person or people, but cannot choose between mutually exclusive candidates for existence.¹⁴

In *Non-identity Case*, you can create either Bob with a good life or Carl with an even better life. Neither option results in comparative or existential harm. Thus, harm-avoidance theories seem to entail that either is permissible. Many now accept this conclusion.¹⁵ However, many others find it counterintuitive; they would claim that Option 1 is impermissible in the *Non-Identity Case*.

Our classification of harm-avoidance theories accommodates both positions. Suppose an agent has options A and B, where A creates some choice-dependent person S, with positive well-being, who would not exist under B, and B creates some choice-dependent person S*, with positive well-being, who would not exist under A. Then A creates *non-identity shortfall* if S’s positive well-being level under A is lower than S*’s well-being level under B. A theory

can recognize both non-identity shortfall and harm as distinct sources of impermissibility. Frick (2020), Otsuka (2017), and Thomas (2022) develop views of this kind, and we do not wish to exclude such views from consideration.

We can classify theories that recognize non-identity shortfall as a source of impermissibility as *harm-avoidance theories* in one of two ways. First, we can recognize non-identity shortfall as a morally significant type of harm, along with comparative and existential harm.¹⁶

The second way, which we favor, involves further restricting the conditions under which the theories that concern us imply that harmless acts are permissible. As stated in §2.1, we are bracketing a list of moral considerations. We can simply add non-identity shortfall to the list. Aside from the *Non-Identity Case* (Table 2), none of the cases that we consider in this paper involve non-identity shortfall; they are *addition* cases, not non-identity cases. Thus, the problems that we raise do not depend on whether non-identity shortfall is treated as a morally relevant consideration.

We can now state the defining feature of harm-avoidance theories. In the domain of normative population ethics, where the considerations relevant to the permissibility of options concern how they affect the number, identities, and well-being levels of people, *and where none of these options create non-identity shortfall*, a harm-avoidance theory is any theory that entails the following commitment:

Harmless Permission: For any option A, if A does not result in comparative or existential harm, then A is permissible.

3 Improvable life avoidance and dominated options

Harmless Permission is compatible with many theories. One simple theory, which often serves as a starting point in discussions of how best to develop a harm-avoidance theory, directs agents to minimize total harm. Suppose the total harm resulting from an option is the sum of all individual harms resulting from it. And suppose the magnitude of an individual

harm is determined as follows. If an option results in existential harm to a person, and there is no alternative under which she has positive well-being, then the magnitude of this harm is just her negative well-being level. Otherwise, it is the difference between her well-being level in the outcome in which she suffers the harm and her well-being level in the outcome in which, among the outcomes the agent can bring about, she is *best off*. Then, according to

Minimize Total Harm: An option A is permissible *iff* there is no alternative B which results in less total harm than A.¹⁷

Minimize Total Harm entails the Asymmetry. Creating a miserable person results in existential harm, and not creating the miserable person results in no harm, all else being equal. Since not creating her results in less total harm than creating her, creating her is impermissible. On the other hand, failing to create a happy person does *no harm*. But nothing could do *less* harm than that. So failing to create a happy person minimizes total harm, and is permissible.

However, as several philosophers have pointed out, Minimize Total Harm faces serious problems.¹⁸ One is The Problem of Improvable Life Avoidance. Another is the problem of requiring dominated options. Table 3 represents a version of a case due to Jacob Ross, which illustrates both problems.¹⁹

	Anne	Bob
Option 1	1	Ω
Option 2	10	10
Option 3	-2000	1000

Table 3: Ross's Case

In *Ross's Case*, Anne suffers 9 units of comparative harm under Option 1, Bob suffers 990 units of comparative harm under Option 2, and Anne suffers 2,010 units of comparative harm under Option 3. Option 1 minimizes total harm. Minimize Total Harm therefore implies that Option 1 is the only permissible option and is morally required.

But this seems implausible. Although Option 3 is clearly impermissible, the claim that Option 2 is impermissible and Option 1 required, is problematic for two reasons.

First, it implies that the agent is required to avoid creating a person with a good life just because this life would be improvable.²⁰ This is The Problem of Improvable Life Avoidance.²¹ If Option 2 is impermissible, this must be because it gives Bob an improvable life—even though this life is *good*. Presumably, if Bob exists, he should prefer a good life to non-existence.

If Bob objects to Option 2 on the grounds that it makes him worse off than he would be under Option 3, we can respond: “the only alternative that avoids even greater harm to someone else would leave you out of existence. Is that what you want?”. One imagines Bob’s answer would be ‘no’. Bob’s objection to Option 2 is what Horton (2021) calls a ‘backfiring objection’.²² The only reasonable way to avoid it involves doing something that, if Bob were to exist, he would, for his own sake, be glad we didn’t do.

Notice that according to Minimize Total Harm, in a choice between only Options 1 and 2, Option 2 minimizes total harm. So in this binary choice Option 2 is required, and hence, permissible, whereas Option 1 is impermissible, and hence, *not* required. But on Minimize Total Harm, Option 2 *becomes* impermissible, and Option 1 *becomes* required, when Option 3 is added to the option set. Minimize Total Harm therefore violates

Improvable Life Acceptance: For any options A and B, if in a binary choice between A and B: (i) A gives person S a good life, (ii) B results in S’s non-existence, and (iii) B is not morally required, then for any option set \mathcal{O} that includes A and B, where (i)–(ii) hold, but S’s life under A is *improvable*, B is not morally required.

Basically, one is not required to leave a person out of existence just to avoid giving her an improvable good life.²³

The second problem with requiring Option 1 is that Option 2 dominates Option 1—it is better for Anne and worse for no one. In fact, Option 2 dominates Option 1 in a very weak

sense. Let us say that Option A *addition-dominates* option B *iff* (i) everyone who exists under B would also exist and be better off under A, (ii) A creates some people who would not exist under B, and (iii) everyone who exists under A has a good life. Next, let us say that A *weakly addition-dominates* B *iff* (iv) A addition-dominates B, and (v) everyone who exists under A has equal well-being.²⁴ Since Minimize Total Harm implies that Option 1 is required in *Ross's Case*, it violates the following plausible principle:

Weak Dominance Addition Exemption: For any options A and B, if A weakly addition-dominates B, then B is not morally required.²⁵

The fact that Minimize Total Harm violates both Improvable-Life Acceptance and Weak Dominance Addition Exemption motivates the search for a harm-avoidance theory that accommodates both.²⁶ In §4, we consider two such theories as case studies before turning to our main thesis (§5) that *any* harm-avoidance theory is deeply problematic.

4 Accommodating improvable lives

Michael McDermott (2019) and Joe Horton (2021) have each proposed a harm-avoidance theory that accommodates both Improvable Life Acceptance and Weak Dominance Addition Exemption. Although their theories differ in important ways, both say that an option is permissible *iff* no one can reasonably object to it, where a necessary condition for someone reasonably objecting to an option is that they would suffer comparative or existential harm under it.

According to McDermott's 'Objection Minimization' theory, a person can reasonably object to an option A *iff* (i) she exists under A, (ii) some alternative B would be better for her (and hence, she suffers harm under A), and (iii) B results in less total harm than A. In *Ross's Case*, Bob cannot reasonably object to Option 2 because Option 3, though better for him, results in more total harm. Anne has no better alternative to Option 2, so she cannot reasonably object

to Option 2 either. Thus, Objection Minimization deems Option 2 permissible.²⁷ This is consistent with Improvable Life Acceptance and Weak Dominance Addition Exemption.

One drawback of Objection Minimization is that it *permits* Option 1.²⁸ Since Option 2 causes more total harm than Option 1, Anne cannot reasonably object to Option 1. Thus, Option 1 is deemed permissible, in violation of the following principle:

Weak Dominance Addition: For any options A and B, if A weakly addition-dominates B, then B is impermissible.

However, permitting weakly addition-dominated options seems less troubling than requiring them, so Objection Minimization is an improvement over Minimize Total Harm.²⁹

Horton's theory—'Avoid Reasonable Objections'—implies that in *Ross's Case* Option 2 is morally required, which seems plausible. Horton's criteria for reasonable objection are complex, and we restate them here with slight adjustments for clarity. These changes do not affect the substance of our argument.³⁰

A person can reasonably object to an option A *iff* she exists, has not consented to A, and there is some alternative option B satisfying 1 through 4.

1. B is better for her than A.
2. B gives a greater sum of well-being than A to the set of people who exist under A.
3. The sum of well-being that B gives to the set of people who exist under B is greater than the sum of well-being that A gives to the set of people who exist under A.³¹
4. No one can reasonably object to B.

Two clarifications are needed. First, Horton thinks that B can be better or worse for someone than A, even if she does not exist under B. For instance, if she has negative well-being under A, but does not exist under B, then on Horton's view, assuming she exists (i.e., A has been

chosen), B is better for her than A. Second, for determining whether condition 2 is satisfied, sums of well-being are assigned to sets of people as follows. The sum of well-being that A gives to the set of people who exist under A is just the sum of well-being for those who exist under A. However, the sum of well-being that B gives to the set of people who exist under A is the sum of the individual well-being values of B for those who exist under A, where some of these people may not exist under B. If someone who exists under A does not exist under B, then, Horton assumes, B gives *zero* well-being to her.

In *Ross's Case*, Avoid Reasonable Objections implies that Option 2 is the only option no one can reasonably object to. Bob cannot reasonably object to Option 2 because the only better option for him, Option 3, gives less well-being to the set {Anne, Bob}. This is sufficient for Option 3 to fail conditions 2 and 3, since Anne and Bob are the only people who exist under Options 2 and 3. Anne cannot reasonably object to Option 2 because no alternative is better for her (condition 1).

Anne can, however, reasonably object to Option 3. Option 2 is better for her (satisfying condition 1), and gives a greater sum of well-being to the set {Anne, Bob}, fulfilling conditions 2 and 3. Since no one can reasonably object to Option 2, condition 4 is also met. Anne can also reasonably object to Option 1. Option 2 is better for her, gives more well-being to the set {Anne}, and to the set {Anne, Bob}, than Option 1, and no one can reasonably object to Option 2. Since Anne can reasonably object to Options 1 and 3, and no one can reasonably object to Option 2, Option 2 is the only permissible option.

Although Avoid Reasonable Objections does not require weakly addition-dominated options, Horton accepts that it sometimes *permits* such options, thus violating Weak Dominance Addition.³²

But Avoid Reasonable Objections and Objection Minimization face another problem, which neither Horton nor McDermott seem to recognize. Both theories violate the following principle:

Non-Sadistic Dominance Addition: For any options A and B, if A weakly

addition-dominates B, and everyone who exists under B has negative well-being, then B is impermissible.

Consider Table 4:

	Ella	Fred
Option 1	-9	Ω
Option 2	1	1
Option 3	-10	1000

Table 4: The Sadistic Dominance Case

In *The Sadistic Dominance Case*, everyone who exists under Option 1 has negative well-being, and Option 2 weakly addition-dominates Option 1. Yet Avoid Reasonable Objections and Objection Minimization imply that Option 1 is *permissible*. It is okay to create only bad lives, when one could instead create only good lives. That is absurdly sadistic.

According to Objection Minimization, in *The Sadistic Dominance Case*, Option 2 results in more total harm than Option 1—999 units of harm for Fred versus 10 for Ella. Moreover, Option 3 is worse for Ella than Option 1. No alternative is better for Ella *and* results in less harm overall than Option 1. So Ella has no reasonable objection to Option 1, making Option 1 permissible.

Avoid Reasonable Objections has the same implication. First, on this theory, no one can reasonably object to Option 3—not even Ella. Option 2 is better for her, but produces less well-being for the set {Ella, Fred}. This violates conditions 2 and 3, since Ella and Fred are the only people who exist under Options 2 and 3. Option 1 is also better for Ella than Option 3, but it produces less well-being for {Ella} than Option 3 does for {Ella, Fred}, thus violating condition 3. Next, since no one can reasonably object to Option 3, Fred can reasonably object to Option 2; he exists under it, does not consent, and Option 3 is better for him and produces more well-being for the set {Ella, Fred}. Finally, since Fred can reasonably object to Option

2, Ella's objection to Option 1 cannot satisfy condition 4. She cannot reasonably object to Option 1, making Option 1 permissible.

This is a serious problem, and it is unclear how to avoid it. Consider how one might try to modify Avoid Reasonable Objections to accommodate Non-Sadistic Dominance Addition. We saw that on Avoid Reasonable Objections, in *The Sadistic Dominance Case*, Fred can reasonably object to Option 2. Now suppose Horton were to modify his theory so that Ella can reasonably object to Option 1, making Option 1 impermissible. Can Fred *still* reasonably object to Option 2? Any answer is problematic.

Suppose Fred *can* reasonably object to Option 2.³³ Then Options 1 and 2 are both reasonably objectionable. Assuming *The Sadistic Dominance Case* is not a moral dilemma, i.e., a case in which every option is impermissible, Option 3 is the only permissible option. But Option 3 is worse for Ella than Option 1. To avoid Ella's objection to Option 1, we are forced to give her an *even worse life* under Option 3. This is an extreme version of what Horton dubbed a "backfiring objection" (§3), which he wants to avoid. It is even more counterintuitive than The Problem of Improvable Life Avoidance. Rather than being required to leave someone out of existence to avoid giving them an improvable good life, we are required to give someone a *more improvable life* to avoid giving them a *less improvable life*.

To see this, notice that according to Avoid Reasonable Objections, Option 3 is not required when the choice is just between Options 1 and 3. According to Harmless Permission, which is the core commitment of any harm-avoidance theory, in such a binary choice, Option 1 is permissible—it causes no comparative or existential harm. But if we then add Option 2 (which is itself impermissible but better for Ella), the theory ends up requiring Option 3, even though Option 3 is worse for Ella than Option 1. This violates a condition that is even more compelling than Improvable Life Acceptance:

Less-Improvable Life Acceptance: For any options A and B, if, in a binary choice between A and B: (i) A results in harm to at most one person S, (ii) S is worse off under A than she is under B, and (iii) B is not morally required, then

for any option set \mathcal{O} that includes A and B, where (i)–(ii) hold, B is not morally required.

To avoid doing what results in harm to a person, one is not thereby required to do what results in even greater harm to that same person.

Condition (i) of the principle, ‘A results in harm to at most one person S’, means A either results in no harm, or results in harm to exactly one person, S. For instance, it is possible that A results in no harm in a binary choice between A and B, but results in comparative harm to S in a larger option set that includes A and B, since, in that larger option set, S’s life under A might be improvable. But it is implausible that this comparative harm under A could *introduce* a moral requirement to choose B, which results in even greater comparative harm to S.

To avoid this, Horton might revise his theory to imply that although Ella can reasonably object to Option 1, Fred *cannot* reasonably object to Option 2. The revised theory would imply that Option 1 is impermissible and Option 2 permissible, which satisfies Non-Sadistic Dominance Addition and Less-Improvable Life Acceptance.

But why can’t Fred reasonably object to Option 2? The harm Fred suffers under Option 2 is greater than the harm Ella suffers under Options 1 and 3. Moreover, Option 3 produces a greater total of well-being than Option 2.

Horton might reply that the harm Ella suffers under Options 1 and 3 is weightier (more important to avoid) than the harm Fred suffers under Option 2, since Ella has negative well-being under Options 1 and 3, whereas Fred has positive well-being under Option 2.

However, we can imagine a variation of *The Sadistic Dominance Case* in which *any number of people* occupy Fred’s position. For instance, suppose that in Table 4 the name ‘Fred’ stands for an enormous group of people (many Freds), each with well-being 1 under Option 2 and well-being 1000 under Option 3. Is Option 2 permissible no matter how many Freds incur 999 units of harm under that Option? If so, then there is no limit to the number of extra people who incur comparative harm under a *permissible* option. Moreover, it seems there is

no limit to *how much* extra comparative harm each of these people can incur. In *The Sadistic Dominance Case*, Fred has well-being 1000 under Option 3; but we can imagine his well-being to be arbitrarily high under that option, and hence, the difference between his well-being under Options 1 and 3 to be arbitrarily large.

Horton's revised theory would then violate the following condition: it is not the case that a permissible option results in *arbitrarily more harm* for each of *arbitrarily many more people* than the alternatives. We can state this condition more carefully as follows. Let P_O be the number of people harmed by option O , and H_O the harm that O imposes on each person. According to

Limit Permissible Harm: For any option set \mathcal{O} , and any option A in \mathcal{O} , there are $i, j > 0$, such that for any option B in \mathcal{O} , if (i) $P_A - P_B > i$, and (ii) $H_A - H_B > j$, then A is impermissible.

This condition is quite weak. It implies that a permissible option cannot *arbitrarily exceed* the alternatives in both the number who suffer harm and the magnitude of that harm. It therefore accommodates a wide range of moral theories, including both prioritarianism, which allows an option to arbitrarily exceed alternatives with respect to the magnitude of comparative harm incurred by a single person, and limited aggregation theories, which allow an option to arbitrarily exceed alternatives with respect to the number of people who incur a *sufficiently small* harm. Rejecting Limit Permissible Harm would make a harm-avoidance theory insufficiently sensitive to comparative harm.

We return to this issue in §5, where we discuss each condition in greater detail.

5 No easy way out

So far, we have considered two harm-avoidance theories. However, as it turns out, *no* harm-avoidance theory can easily preserve all three conditions introduced in §4—Non-Sadistic

Dominance Addition, Less-Improvable Life Acceptance, and Limit Permissible Harm. Choice contexts like *The Sadistic Dominance Case* reveal a tension between these conditions and the core commitment of harm-avoidance theories, Harmless Permission. The only way to preserve Harmless Permission without violating one of these conditions is to violate a further condition we call *No Dilemmas*, which states that choice contexts like *The Sadistic Dominance Case* are *not* moral dilemmas—contexts in which every option is impermissible.

We prove this result formally in §6; but the idea can be grasped by looking at Table 5, which represents the modified version of *The Sadistic Dominance Case* considered in §4.

	Ella	n People
Option 1	−9	Ω
Option 2	1	1
Option 3	−10	1000

Table 5: Modified Sadistic Dominance Case

Suppose the number of people, n , who exist under Options 2 and 3 is sufficiently large, and the difference between their well-being levels under Options 2 and 3 sufficiently great, that Option 2 is impermissible according to Limit Permissible Harm.

Now in a choice between only Options 1 and 3, Option 1 does not result in comparative or existential harm; so, according to Harmless Permission, in that binary choice, Option 1 is permissible, and Option 3 is not required. Hence, according to Less-Improvable Life Acceptance, Option 3 is not required in a choice between all three options. (When Option 2 is added alongside Options 1 and 3, Ella is the only person harmed under Option 1, but she is worse off under Option 3 than under Option 1.) Next, if Option 3 is not required in a choice between all three options, i.e., Option 3 is not the only permissible option in that context, and No Dilemmas is true, i.e., at least one option is permissible, then *either Option 1 or Option 2 is permissible*. But according to Non-Sadistic Dominance Addition, Option 1 is not permissible,

as it gives everyone who exists negative well-being and is weakly addition-dominated by Option 2; and, according to Limit Permissible Harm, Option 2 is not permissible. So *it is not the case that* either Option 1 or Option 2 is permissible. Contradiction.

Those who accept Harmless Permission must either reject Non-Sadistic Dominance Addition, Less-Improvable Life Acceptance, Limit Permissible Harm, or No Dilemmas.

How should harm-avoidance theorists respond? We briefly consider the possibilities in the remainder of this section.

5.1 Rejecting No Dilemmas

In our view, the least promising response is to treat *The Modified Sadistic Dominance Case* as a moral dilemma, and hence, reject No Dilemmas.

First, note that we have defined a ‘moral dilemma’ as a context in which every option is impermissible. This definition is somewhat general. Typically, a ‘moral dilemma’ is defined as a context in which an agent will act impermissibly no matter she does *because* each option is *morally required* but she cannot choose all options.³⁴

The possibility of moral dilemmas highly contested. But beyond this, treating *The Modified Sadistic Dominance Case* as a moral dilemma seems *ad hoc*. We cannot see any independent reason to view this case as a genuine moral dilemma, as opposed to a hard *moral conflict*—a situation involving strong but competing moral reasons.

A genuine moral dilemma is supposed to feel *irresolvable* in the sense that none of the alleged conflicting requirements overrides the others.³⁵ But *The Modified Sadistic Dominance Case* does not seem irresolvable. For example, it seems to us that the agent’s strongest moral reason is to avoid Option 1, which results in negative well-being for everyone who exists. Some readers may instead think the strongest reason is to avoid Option 2 or 3. But that judgment also cuts against the claim that the conflict is irresolvable.

5.2 Rejecting Non-Sadistic Dominance Addition

The second-least promising response in our view is to reject Non-Sadistic Dominance Addition. Rejecting this condition is not unprecedented; critical-level utilitarianism violates it, for instance.³⁶ But critical-level utilitarianism implies what Gustaf Arrhenius (2000, p. 63) calls the *Sadistic Conclusion*: for any number of lives at any negative well-being level (horribly tormented lives) there are choice situations in which it is better to create these lives rather than some number of lives with positive well-being. Rejecting Non-Sadistic Dominance Addition might open the door to some version of this conclusion, for instance, one according to which it is *permissible* to create horribly tormented lives when one could create only good lives. Even if we could somehow avoid that extreme conclusion, giving up Non-Sadistic Dominance Addition is a significant cost.

5.3 Rejecting Less-Improvable Life Acceptance

Here, we consider one possible rationale for rejecting Less-Improvable Life Acceptance.

Harmless Permission implies there is no requirement to create people with positive well-being. This explains why, in the *Modified Sadistic Dominance Case*, Option 3 is not required when it is the only alternative to Option 1. Yet it seems bizarre that adding Option 2 could *make* Option 3 required. When Option 2 is added, Option 1 results in comparative harm to Ella, since she is worse off under Option 1 than under Option 2. But Ella is even worse off under Option 3 than under Option 1. So Option 3 is morally worse than Option 1 *in this respect*. The only other difference between the two options is that Option 3 creates n people with well-being 1,000. If there is no requirement to create people with positive well-being, where does the requirement to choose Option 3 come from?

Here is one possible explanation. Following Theron Pummer (2023, §1.1), suppose there are two kinds of moral reasons. A *requiring reason* to do an act A serves to make A required. By contrast, a *permitting reason* to do an act A prevents the balance of requiring reasons from

making A wrong but cannot *by itself* generate a moral requirement to perform A. An act B is in fact required (wrong not to do) *iff* there is most requiring reason overall to do B and no sufficiently strong permitting reason not to do B.³⁷

But suppose, *pace* Pummer (2023), permitting reasons can create moral requirements *indirectly* in the following way. If one has a requiring reason to avoid option A, and a (possibly stronger) requiring reason to avoid option B, then one is required to choose B over A if there is also a sufficiently strong permitting reason to choose B, but no such reason to choose A, and all else is equal. The permitting reason to choose B diffuses, or neutralizes, the strength of one's requiring reason to choose A, thereby tipping the balance of one's requiring reasons in favor of B. Call this view *indirect requirements*.³⁸

Now suppose we have requiring reasons to avoid harm, but permitting reasons to create people with good lives.³⁹ In *The Modified Sadistic Dominance Case*, there is a requiring reason against Option 1, since it results in harm to Ella, and an even stronger requiring reason against Option 3, since it results in even greater harm to Ella. But if we choose Option 1, the requiring reason we fail to comply with is not neutralized by any permitting reason, since Option 1 does not create people with good lives. However, the requiring reason we fail to comply with by choosing Option 3 is, perhaps, neutralized by the permitting reason to create *n* people with good lives. Hence, according indirect requirements, Option 3 is required in a choice between Options 1–3. This, however, does *not* imply that Option 3 is required in a choice between only Options 1 and 3.

But indirect requirements is not only a controversial view of permitting (or justifying) reasons; it is also incompatible with the Asymmetry. Consider Table 6:

	Georgia	Hank	Irving
Option 1	100	Ω	Ω
Option 2	101	-1000	Ω
Option 3	100	Ω	1000

Table 6: The All Else Equal Case

What should we say about Options 1 and 3? There are equally strong requiring reasons to avoid Options 1 and 3, since each results in comparative harm of 1 to Georgia. But there is a permitting reason to choose Option 3 (and none to choose Option 1), since Option 3 creates Irving with a good life. So according to indirect requirements, if all else is equal, we are morally required to choose Option 3.

And all else is equal. Options 1 and 3 give Georgia the same well-being level, and both avoid the existential harm to Hank, which he would suffer under Option 2. The only difference is Option 3 creates Irving with positive well-being, and Option 1 creates no one. Indirect requirements therefore implies that one can be required to create people with good lives rather than not create them, all else being equal.

But according to the Asymmetry there is *no* moral requirement to create people with good lives rather than not create them, all else being equal. Indirect requirements is incompatible with the Asymmetry.

Perhaps harm-avoidance theorists can respond by formulating the Asymmetry so that it applies only to choice contexts with *exactly two* options—creating a person and not creating them. That would make indirect requirements consistent with the letter, but not the spirit, of the Asymmetry. For again, in the *All Else Equal Case*, the only consideration for Option 3, in comparison with Option 1, is that it creates someone with positive well-being. If we are not required to create such a person when we have only two options, how could we be required to do so in this case?⁴⁰

We think harm-avoidance theorists should reject indirect requirements. But then they cannot appeal to it as a justification for rejecting Less-Improvable Life Acceptance.

5.4 Rejecting Limit Permissible Harm

The last response is to reject Limit Permissible Harm. To motivate this response, one could claim that preventing harm to choice-independent people has absolute priority over preventing comparative harm to choice-dependent people, provided that the latter would have good lives.⁴¹ This entails that in *The Modified Sadistic Dominance Case* Option 2 is required, no matter how much comparative harm n choice-dependent people incur under that option.⁴²

To support this claim, one could argue that when a choice-dependent person suffers comparative harm, but has a good life, she does not suffer a *genuine harm*. There is even a sense in which she *benefits*, by coming into existence with a good life.⁴³

However, even if a comparative harm to a well-off choice-dependent person is not a genuine harm, it is a morally significant welfare shortfall. Suppose I can either (i) create someone with a fantastic life, (ii) create her with a life barely worth living, or (iii) create no one. Option (ii) seems impermissible. I cannot reasonably defend the choice of (ii) by pointing out that the person I created has a good life and would not exist under (iii). I have moral reason not to make a choice-dependent person worse off than I could have made her.

Moreover, the claim that avoiding harm to a choice-independent person has absolute priority over avoiding any comparative harm to any number of well-off choice-dependent people is extreme. Consider Table 7, which represents a version of Derek Parfit's *Mere Addition Paradox*:⁴⁴

	Josie	n People
Option 1	100	Ω
Option 2	101	1
Option 3	99	99

Table 7: The Mere Addition Case

Josie is the only choice-independent person, and n the number of choice-dependent people under Options 2 and 3. On the view we are now considering, Option 2 is required, as any alternative results in harm to Josie. But our strong intuition is that Option 2 is *impermissible* even for $n = 1$. This intuition is stronger for larger values of n . We expect most readers, even most harm-avoidance theorists, to share this intuition.⁴⁵

Rejecting Limit Permissible Harm is a high cost, even if it is less costly than rejecting one of the other conditions.

6 An impossibility theorem

In §5, we gave an informal argument for our main result. In this section, we prove this result.

Consider Table 8.

	Person 1	Person 2	...	Person n
Option 1	$-x$	Ω	...	Ω
Option 2	y	y	...	y
Option 3	$-x - \varepsilon$	z	...	z

Table 8: The General Schema

The General Schema represents a wide range of possible choice-contexts where x , y , z , and ε are possible well-being values for Persons 1 ... n , with $n > 2$.⁴⁶ As in previous cases, Ω

represents non-existence for a person.

The schema has four important features:

Feature 1: For any $x, y, z, \varepsilon > 0$, in a binary choice between Option 1 and Option 3, Option 1 results in no existential and no comparative harm.

Feature 2: For any $x, y, z, \varepsilon > 0$,

- (a) In a choice between all three options, Option 1 results in harm only to Person 1.
- (b) In a binary choice between Option 1 and Option 3, and in a choice between all three options, Person 1 is worse off under Option 3 than Option 1.

Feature 3: For any $x, y, z, \varepsilon > 0$, if $z - y > y - (x - \varepsilon)$, then in a choice between all three options, Option 2 harms $n - 1$ more people than Options 1 and 3, and each of these people incurs $(z - y) - (y - (-x - \varepsilon))$ more harm than the greatest harm suffered by Person 1 under her worst option (Option 1).

Feature 4: For any $x, y > 0$,

- (a) Option 2 weakly addition-dominates Option 1.
- (b) Everyone who exists under Option 1 has negative well-being.

We make two assumptions:

Deontic Bivalence: An option is impermissible *iff* it is not permissible.

Moral Requirement: An option is morally required *iff* it is permissible and any alternative is impermissible.

Given Deontic Bivalence and Moral Requirement, we demonstrate that Harmless Permission, Non-Sadistic Dominance Addition, Less-Improvable Life Acceptance, Limit Permissible Harm, and No Dilemmas are jointly incompatible.

Proof. Assume for reductio

(1) Harmless Permission: For any option A, if A does not result in comparative or existential harm, then A is permissible.

(2) Less-Improvable Life Acceptance: For any options A and B, if, in a binary choice between A and B: (i) A results in harm to at most one person S, (ii) S is worse off under A than she is under B, and (iii) B is not morally required, then for any option set \mathcal{O} that includes A and B, where (i)–(ii) hold, B is not morally required.

(3) Non-Sadistic Dominance Addition: For any options A and B, if A weakly addition-dominates B, and everyone who exists under B has negative well-being, then B is impermissible.

(4) Limit Permissible Harm: For any option set \mathcal{O} , and any option A in \mathcal{O} , there are $i, j > 0$, such that for any option B in \mathcal{O} , if (i) $P_A - P_B > i$, and (ii) $H_A - H_B > j$, then A is impermissible.

(5) No Dilemmas: For any $x, y, z, \varepsilon > 0$, at least one of Options 1 through 3 is permissible.

From (1), Feature 1, Moral Requirement, and Deontic Bivalence,

(6) For any $x, y, z, \varepsilon > 0$, in a binary choice between Option 1 and Option 3, Option 3 is not morally required.

From (2), (6), and Feature 2,

(7) For any $x, y, z, \varepsilon > 0$, in a choice between all three options, Option 3 is not morally required.

From (5), (7), Moral Requirement, and Deontic Bivalence,

(8) For any $x, y, z, \varepsilon > 0$, in a choice between all three options, either Option 1 is permissible or Option 2 is permissible.⁴⁷

From (4), Deontic Bivalence, and Feature 3,

(9) For some $x, y, z, \varepsilon > 0$, $n - 1 > i$, and $(z - y) - (y - (-x - \varepsilon)) > j$, in a choice between all three options, Option 2 is impermissible.

From (8), (9), and Deontic Bivalence,

(10) For some $x, y, z, \varepsilon > 0$, $n - 1 > i$, and $(z - y) - (y - (-x - \varepsilon)) > j$, in a choice between all three options, Option 1 is permissible.

But from (3) and Feature 4,

(11) For any $x, y, z, \varepsilon, > 0$, in a choice between all three options, Option 1 is impermissible.

So, from (11) and Deontic Bivalence,

C. It is not the case that for some $x, y, z, \varepsilon > 0$, $n - 1 > i$, and $(z - y) - (y - (-x - \varepsilon)) > j$, in a choice between all three options, Option 1 is permissible.

□

C contradicts (10). We must therefore reject at least one of Harmless Permission, Non-Sadistic Dominance Addition, Less-Improvable Life Acceptance, Limit Permissible Harm, and No Dilemmas.

7 Conclusion

Harm-avoidance theories provide one way of grounding the Asymmetry. A challenge for any harm-avoidance theory is providing an adequate response to The Problem of Improvable Life

Avoidance and the problem of requiring weakly addition-dominated options. Some harm-avoidance theories can meet this challenge. But we have shown that any harm-avoidance theory, any theory committed to Harmless Permission, either permits weakly addition-dominated options in which everyone has negative well-being, requires giving a person a more improvable life just to avoid giving that same person a less improvable life, permits allowing any amount of extra comparative harm to befall any number of people, or accepts moral dilemmas. The prospects for harm-avoidance theories seem bleak.

What does this mean for the Asymmetry? Those drawn to it may want to reject Harmless Permission in favor of an alternative, Asymmetry-friendly principle. Indeed, there may be independent reason to reject Harmless Permission. One limitation of this principle is that it treats only comparative and existential harm as morally significant. Yet, as Elizabeth Harman (2009) and Seana Shiffrin (1999) have argued, there is also what we can call *intrinsic harm*—being in a state with intrinsically bad aspects such as pain or bodily damage. Unlike comparative or existential harm, intrinsic harm does not presuppose that one could have been better off under some alternative or that there exists an alternative in which one does not exist.

Recognizing intrinsic harm captures the intuition that someone with a miserable life suffers harm even if she would exist under every alternative. For instance, it lets us say that in the *Sadistic Dominance Case*, Ella suffers harm under Option 1, even when the choice is only between Options 1 and 3, either of which would give Ella negative well-being.

But if Harmless Permission is rejected on these grounds, what should replace it? We think any alternative Asymmetry-friendly principle will still be incompatible with the four principles we identified—Non-Sadistic Dominance Addition, Less-Improvable Life Acceptance, Limit Permissible Harm, and No Dilemmas. The source of the tension between these principles and Harmless Permission is that the latter implies that in a simple binary choice, an agent is not required to make an existing person with negative well-being worse off in order to create new people with positive well-being. Any principle with this implication will face

the same challenge as harm-avoidance theories; it will be incompatible with the other four principles.

Yet this implication seems unavoidable for any principle that respects the Asymmetry. If, as the Asymmetry holds, we are not required to create someone with positive well-being, then it is hard to see how we could be required to make someone worse off in order to do so.⁴⁸

Thus our challenge to harm-avoidance theories extends to the Asymmetry itself: the task is to find a general theory that both supports the Asymmetry and is compatible with the four principles. That looks to be a tall order.

Notes

¹Inspired by Thomas (2022, p. 471), we interpret ‘all else being equal’ to mean setting aside affects of the agent’s choice on those who exist independently of that choice.

²(Chappell, 2017; Roberts, 2011; Thomas, 2022; Thornley, 2023). The Asymmetry has also been couched in terms of reasons (Bader, 2022; Frick, 2020; McMahan, 1981).

³(McDermott, 2019, p. 437; Thomas, 2022, p. 473). See also Melinda Roberts (2011) and Larry Temkin (2012).

⁴For an extensive defense of this view, see David Boonin (2014), especially chapters 4 and 5.

⁵(Arrhenius et al., 2022, p. 1)

⁶On this distinction, see (Bykvist, 2006; McMahan, 1981, 2013; Parfit, 2017; Podgorski, 2023; Thomas, 2022). Some argue that a person can suffer harm in a different sense, that of being in a state with an intrinsically bad aspect (Harman, 2009; Shiffrin, 1999). We say more about this in §7.

⁷(Thomas, 2022, §4.1)

⁸Representing well-being numerically in this way assumes interpersonal well-being comparability on a ratio scale—a strong but standard assumption in the literature.

⁹The term ‘improvable life’ is originally from Ross (2015, p. 443).

¹⁰(Horton, 2021, pp. 486–7; Podgorski, 2023, p. 345 states his theory in terms of complaints; see also Pummer, 2024, footnote 5).

¹¹For instance, Boonin (2014) adopts the terminology of harm, but Otsuka (2017) the terminology of com-

plaints.

¹²See Part IV of *Reasons and Persons* (Parfit, 1984).

¹³(Arrhenius, 2003, footnote 4)

¹⁴For example, she might choose between adding Bob alone and adding Bob together with Carl, but she cannot choose between adding Bob without Carl and adding Carl without Bob.

¹⁵(Boonin, 2014; Horton, 2021; McDermott, 2019; Podgorski, 2023; Roberts, 2007; Spencer, 2021)

¹⁶See, e.g., Meacham (2012).

¹⁷(McDermott, 2019, p. 440)

¹⁸See Thomas (2022, §2) for a full discussion of the difficulties with this theory.

¹⁹(Ross, 2015, p. 440). Our presentation of the case was sourced, with minor changes, from Podgorski (2023, p. 353).

²⁰(Horton, 2021; McDermott, 2019; Podgorski, 2023)

²¹Ross (2015, p. 443) characterizes “The fundamental problem,” as that we have “reason to prefer outcomes in which a given person does not exist to outcomes in which this person exists and has an improvable life.”

²²See also McDermott (2019).

²³Podgorski expresses a similar intuition: “It should not be possible to start with a set of choices which permit us to create someone with a happy life, add an option under which they are better off, and thereby generate a complaint *on their behalf* which makes it impermissible to create them at all” (2023, p. 354).

²⁴(cf. Thornley, 2023, pp. 522–3)

²⁵This formulation is inspired by Elliot Thornley, who proposes a similar principle in his criticism of a theory defended by Horton, which we consider in §4. Thornley (2023, p. 522) calls his principle ‘Weak Normative Dominance Addition’, which is like ours except it says that if everyone has non-negative well-being in the dominated option, then if that option is permissible, the dominating option is permissible.

²⁶For instance, Horton (2021), McDermott (2019), and Podgorski (2023).

²⁷(McDermott, 2019, p. 444)

²⁸Thomas (2022, footnote 8) points this out.

²⁹(cf. Thomas, 2022, footnote 23)

³⁰For his original statement, see (Horton, 2021, p. 499).

³¹We here include Thornley’s (2023, p. 519) amendment to Horton’s condition 3.

³²He also accepts (2021, 496–7, §4.2) that it sometimes implies that a weakly-addition dominated option is permissible while the option that weakly addition-dominates it is impermissible. This objection to Avoid Reasonable Objections is emphasized by Thornley (2023, pp. 521–2).

³³Note that this violates condition 4—Ella can reasonably object to Option 1 even though the only better option for her (Option 2) is also reasonably objectionable.

³⁴(McConnell, 2024)

³⁵(Sinnott-Armstrong, 1988, Ch. 1)

³⁶See, e.g., Blackorby et al. (2005). Here we construe critical-level utilitarianism as both an axiological and a deontic theory, one that provides both a value ordering of outcomes and a moral principle telling agents to choose the best outcome.

³⁷Other frameworks have been proposed along the same lines as Pummer’s, such as (Gert, 2004; Kaczmarek & Lloyd, 2025; Kamm, 1996; Lazar, 2013; Muñoz, 2021; Pummer, 2023). See Little and Macnamara (2021) for an overview of the literature on so-called ‘non-requiring reasons’. We adopt Pummer’s framework because it is the simplest for our purposes.

³⁸This view departs from the standard position of proponents of permitting (or purely justifying) reasons. Orthodoxy has it that such reasons directly alter deontic status, without upsetting the balance of the all-things-considered requiring strength of one’s reasons (Kaczmarek & Lloyd, 2025, p. 5). For instance, Pummer himself claims that while permitting reasons can remove requirements, they cannot create them, even indirectly. Thomas (2022) also seems committed to rejecting this possibility.

³⁹To our knowledge, Thomas (2022, §5) is the only philosopher to formulate a fully general theory incorporating permitting (or purely justifying) reasons to create happy people. However, many have discussed similar ideas in this domain, including Arrhenius (2009), McMahan (2013), Mogensen (2019), Podgorski (2023), and Roberts (2011).

⁴⁰[redacted] develops this argument.

⁴¹Without the qualification that they have good lives, the claim is absurdly sadistic. It implies that we should avoid any amount of harm to a choice-independent person by creating any number of choice-dependent people with negative well-being.

⁴²Michael Otsuka (2017) suggests a view like this. He writes: “When the complaints of those whose existence is choice-independent come into conflict with the complaints of those whose existence is choice-dependent, the complaints of the former trump the complaints of the latter” (2017, 212, footnote 18).

⁴³We thank [redacted] for this point.

⁴⁴(1984, pp. 425–433)

⁴⁵For instance, Horton’s, McDermott’s, and Thomas’s theories all imply that Option 2 is impermissible.

⁴⁶Imposing this restriction on n is necessary to make Limit Permissible Harm relevant; otherwise the number of people harmed under an option cannot exceed the number harmed under any alternative.

⁴⁷Omitting the quantifier ‘For any $x, y, z, \varepsilon > 0$ ’ for convenience, the derivation proceeds as follows. From

Moral Requirement, Option 3 is morally required *iff* the following conjunction is true: *Option 3 is permissible and Option 1 is impermissible and Option 2 is impermissible*. According to (7), Option 3 is not morally required, so the conjunction is false. From the negation of the conjunction, Deontic Bivalence and De Morgan's Law, we derive the following disjunction: *either Option 3 is not permissible or Option 1 is permissible or Option 2 is permissible*. If Option 3 is permissible, then the first disjunct is false, and hence, by Disjunctive Syllogism, we derive (8): *either Option 1 is permissible or Option 2 is permissible*. If Option 3 is not permissible, then by No Dilemmas, which says that at least one of the three options is permissible, we also derive (8): *either Option 1 is permissible or Option 2 is permissible*.

⁴⁸See [redacted] for a defense of this claim.

References

- Arrhenius, G. (2000). *Population ethics: The challenge of future generations* [Doctoral dissertation, Uppsala University].
- Arrhenius, G. (2003). The person-affecting restriction, comparativism, and the moral status of potential people. *Ethical Perspectives*, 10(3), 185–195.
- Arrhenius, G. (2009). Can the person-affecting restriction solve the problems in population ethics? In M. Roberts & D. Wasserman (Eds.), *Harming future persons: Ethics, genetics and the nonidentity problem* (pp. 289–314). Springer.
- Arrhenius, G., Bykvist, K., Campbell, T., & Finneron-Burns, E. (2022). Introduction. In G. Arrhenius, K. Bykvist, T. Campbell, & E. Finneron-Burns (Eds.), *The Oxford handbook of population ethics* (pp. 1–12). Oxford University Press.
- Bader, R. (2022). The asymmetry. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (Eds.), *Ethics and existence: The legacy of Derek Parfit* (pp. 15–37). Oxford University Press.
- Blackorby, C., Bossert, W., & Donaldson, D. J. (2005). *Population issues in social choice theory, welfare economics, and ethics*. Cambridge University Press.

- Boonin, D. (2014). *The non-identity problem & the ethics of future people*. Oxford University Press.
- Bykvist, K. (2006). The benefits of coming into existence. *Philosophical Studies*, 135(3), 336–362.
- Chappell, R. Y. (2017). Rethinking the asymmetry. *Canadian Journal of Philosophy*, 47(2-3), 167–177.
- Frick, J. (2020). Conditional reasons and the procreation asymmetry. *Philosophical Perspectives*, 34(1), 53–87.
- Gert, J. (2004). *Brute rationality*. Cambridge University Press.
- Harman, E. (2009). Harming as causing harm. In M. Roberts & D. Wasserman (Eds.), *Harming future persons: Ethics, genetics and the nonidentity problem* (pp. 137–154). Springer.
- Horton, J. (2021). New and improvable lives. *Journal of Philosophy*, 118(9), 486–503.
- Kaczmarek, P., & Lloyd, H. R. (2025). Moral uncertainty, pure justifiers, and agent-centred options. *Australasian Journal of Philosophy*, 00, 1–21. <https://doi.org/10.1080/00048402.2025.2508226>.
- Kamm, F. M. (1996). *Morality, mortality, vol. 2, rights, duties, and status*. Oxford University Press.
- Lazar, S. (2013). Associative duties and the ethics of killing in war. *Journal of Practical Ethics*, 1(1), 3–48.
- Little, M. O., & Macnamara, C. (2021). Non-requiring reasons. In R. Chang & K. Sylvan (Eds.), *The Routledge handbook of practical reasons* (pp. 393–404). Routledge.
- McConnell, T. (2024). Moral dilemmas. In E. N. Zalta & U. Nodelman (Eds.), *The stanford encyclopaedia of philosophy (spring 2024 edition)*.
- McDermott, M. (2019). Harms and objections. *Analysis*, 79(3), 436–448.
- McMahan, J. (1981). Problems of population theory. *Ethics*, 92(1), 96–127.
- McMahan, J. (2013). Causing people to exist and saving lives. *Journal of Ethics*, 17(1/2), 5–35.

- Meacham, C. J. (2012). Person-affecting views and saturating counterpart relations. *Philosophical Studies*, 158(2), 257–287.
- Mogensen, A. (2019). Staking our future: Deontic long-termism and the non-identity problem. *Global Priorities Institute Working Paper No. 9-2019, 00*, 1–32. <https://globalprioritiesinstitute.org/andreas-mogensen-staking-our-future-deontic-long-termism-and-the-non-identity-problem/>.
- Muñoz, D. (2021). From rights to prerogatives. *Philosophy and Phenomenological Research*, 102(3), 608–623.
- Otsuka, M. (2017). How it makes a moral difference that one is worse off than one could have been. *Politics, Philosophy & Economics*, 17(2), 192–215.
- Parfit, D. (1984). *Reasons and persons*. Oxford University Press.
- Parfit, D. (2017). Future people, the non-identity problem, and person-affecting principles. *Philosophy & Public Affairs*, 45(2), 118–157.
- Podgorski, A. (2023). Complaints and tournament population ethics. *Philosophy and Phenomenological Research*, 106(2), 344–367.
- Pummer, T. (2023). *The rules of rescue: Cost, distance, and effective altruism*. Oxford University Press.
- Pummer, T. (2024). Future suffering and the non-identity problem. *Global Priorities Institute Working Paper No. 17-2024, 00*, 1–50. <https://globalprioritiesinstitute.org/future-suffering-and-the-non-identity-problem-theron-pummer/>.
- Roberts, M. (2007). The non-identity fallacy: Harm, probability and another look at Parfit's depletion example. *Utilitas*, 19(3), 267–311.
- Roberts, M. (2011). The asymmetry: A solution. *Theoria*, 77(4), 333–367.
- 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.

- Sinnott-Armstrong, W. (1988). *Moral dilemmas*. Basil Blackwell.
- Spencer, J. (2021). The procreative asymmetry and the impossibility of elusive permission. *Philosophical Studies*, 178(11), 3819–3842.
- Temkin, L. (2012). *Rethinking the good*. Oxford University Press.
- Thomas, T. (2022). The asymmetry, uncertainty, and the long term. *Philosophy and Phenomenological Research*, 107(2), 470–500.
- Thornley, E. (2023). The procreation asymmetry, improvable-life avoidance and impairable-life acceptance. *Analysis*, 83(3), 517–526.

The Institute for Futures Studies is an independent research foundation financed by contributions from the Swedish Government and through external research grants. The institute conducts interdisciplinary research on future issues and acts as a forum for a public debate on the future through publications, seminars and conferences.

This working paper has been produced by the Mimir Center for Long Term Futures Research at the Institute for Futures Studies.

© Institute for Futures Studies and the author