

Tim Campbell ¹ & Patrick Kaczmarek²

Improving Lives and Avoiding Harm: A Critical Response to Harm-Based Arguments for Climate Anti-Natalism³

Creating a new person produces more CO2 emissions than many other lifestyle choices, such as driving a gas-powered car, eating meat, and flying. According to Climate Anti-Natalism, for this reason, in many instances, it is wrong to create a new person, even if that person would have a good life. Arguments for Climate Anti-Natalism point to the harm that CO2 emissions cause, but they do not recognize any moral reason to create people with good lives. We identify a harm-avoidance principle underlying arguments for Climate Anti-Natalism. We then show that any moral theory that accommodates this harm-avoidance principle has implausible implications. Such a theory either permits agents to create people with bad lives rather than with good lives, requires agents to harm people just to avoid imposing less harm on those same people, or permits agents to impose any amount of uncompensated harm. A reasonable response to this problem is to reject the harm-avoidance principle, thus undermining the case for Climate Anti-Natalism.

¹ Institute for Futures Studies, Stockholm and Mimir Center for Long Term Futures Research, Stockholm, timothy.campbell@iffs.se.

² Centre for the Study of Existential Risk, Cambridge, pakazmarek@gmail.com.

³ Funding from Effective Giving and Riksbankens Jubileumsfond (grant number: P22-0662) is gratefully acknowledged.

1. Introduction

Empirical evidence suggests having a child generates more CO₂ emissions than driving a gas-powered car, flying, and eating meat (Wynes and Nicholas 2017). At the same time, several ethicists have argued that producing large quantities of CO₂ by having children is, morally speaking, no different from producing comparable quantities of CO₂ by such other means (Young 2001; MacIver 2015; Conly 2015; Hedberg 2019, Rieder 2018; Burkett 2021). Some have gone as far as to claim that for many people in rich countries, having children is morally wrong, since the CO₂ emissions from having children contribute to the harm of climate change and are inessential for a decent life (Burkett 2021). Let us call the claim that it is wrong to have children because it would contribute to the harm of climate change *climate anti-natalism*.

The arguments for climate anti-natalism neglect some important questions. This paper focuses on the most central neglected question: can the fact that a person would exist with a good life, by itself, provide a moral reason to create this person? Once this question is brought into focus, it quickly becomes apparent that climate anti-natalists have no easy answer to it. For grappling with the question requires us to dive into the thorny field of normative population ethics, in which every theory is deeply problematic. The theories most naturally suited to climate anti-natalism are those built around the idea that while we have moral obligations to avoid harming people, we have either no moral reason whatsoever, or at least no requiring moral reason, to create people with good lives.

The debate around climate anti-natalism is therefore related to debates in population ethics about what has become known as *the Asymmetry*. In its *moral reasons* variation, the *Asymmetry* states that, when other things are equal, we have moral reason to avoid creating people with bad lives, but no moral reason to create people with good lives when the alternative is to create no one. The *Asymmetry* also has a deontic formulation, according to which, when other things are equal, we are morally required not to create people with bad lives, but not to create people with good lives when the alternative is to create no one. Underlying the deontic version of the *Asymmetry* is a claim that we call *harm-avoidance*: it can be morally impermissible to harm individuals, but refraining from creating individuals with good lives is morally permissible, other things being equal. More generally, some have claimed that in a certain restricted class of choice situations, which we identify in §2.1, an option is impermissible only if it does harm. Call this *the harm-avoidance account* of the *Asymmetry*.

Several moral theories are based on the harm-avoidance account. Call any such theory a *harm-avoidance theory*. As many have acknowledged, the simplest harm-avoidance theories, those which morally require agents to minimize total harm, face an

especially troubling problem, *the Problem of Improvable Life Avoidance*, which we present in §3.

The search for a harm-avoidance theory that deals adequately with *the Problem of Improvable Life Avoidance* is underway. In this paper, we argue that this search is ill-starred. Any harm-avoidance theory either faces some no less troubling variant of *the Problem of Improvable Life Avoidance*, permits agents to create people with miserable lives rather than with good lives, or permits inflicting any amount of harm, thus giving inadequate consideration to harm-avoidance.

We begin, in §2, by making ‘harm-avoidance’ more precise, explaining in greater detail what we take to count as a harm-avoidance theory, and identifying a defining feature of any such theory, a commitment that we call *Harmless Permission* (cf. §2.4). In §3, we present *the Problem of Improvable Life Avoidance*. In §§4–5, we scrutinize four state-of-the-art harm-avoidance theories as case studies, specifically those recently put forward by Michael McDermott, Joe Horton, Teruji Thomas, and Abelard Podgorski. These theories seem to offer an adequate response to *the Problem of Improvable Life Avoidance*, but they encounter other, no less severe, problems. By exposing the problems these theories face, we identify three plausible principles: *Weaker Dominance Addition* (cf. §4), *Weak Improvable Life Acceptance* (cf. §5.1), and *Limit Permissible Harm* (cf. §5.2). In §6, we prove that no harm-avoidance theory can satisfy all three principles. Finally, in §7, we conclude with some reflections on what this means for climate anti-natalism. Our result casts doubt on the harm-avoidance account, and points toward the existence of moral requirements to create people with good lives. In other words, our result undermines the case for climate anti-natalism. While it is still possible to defend climate anti-natalism on grounds other than harm-avoidance, a convincing alternative case is yet to be made.

2. The Harm-Avoidance Account and General Theories

According to the harm-avoidance account, in certain choice situations, an option is impermissible only if it does harm. In this section, we spell out what the relevant choice situations are, and what counts as a harm-avoidance theory.

There are four points of clarification regarding harm-avoidance theories: their domain of application (§2.1); their characterization of ‘harm’ (§2.2); the various conceptual framings they can adopt (§2.3); and what they take to be the moral significance of creating well-off people (§2.4). Regarding the final point, according to our classification, a harm-avoidance theory can recognize a special exception to the rule that harmless options are permissible, namely an exception for *non-identity cases*, or

cases in which the agent can create a new person (or people) but can determine which new person (or people) will exist.

2.1 Domain restriction

A theory based on the Asymmetry *could* be a fully general theory of the permissibility of options.⁴ However, most theories we will discuss are not clearly intended to be fully general. Their proponents either explicitly assume, or are most charitably interpreted as assuming, that the application of the theories they defend is restricted to *normative population ethics*—the part of moral philosophy concerned with the permissibility of options that may affect the (i) number, (ii) identities, and (iii) well-being levels of people, and where permissibility facts supervene on facts about (i)–(iii). Considerations other than (i)–(iii), such as personal virtue, agent-relative prerogatives, special obligations, and whether agents lie, cheat, and steal (without affecting the number, identities, and well-being levels of people) are typically sidelined.⁵ Among the theories we discuss in §4–5, there is one minor exception. Joe Horton (2021) defends a theory, which we shall include in our classification of harm-avoidance theories, according to which a necessary condition for an option to be impermissible is that it affects a *non-consenting* individual. But this minor exception aside, harm-avoidance theories are concerned only with how considerations (i)–(iii) affect permissibility.

2.2 Harm

The second point of clarification concerns ‘harm’. Harm-avoidance theories recognize only two types of harm as morally significant: comparative harm and existential harm.⁶

Suppose your only options are A and B. If a certain person exists given the choice of either A or B, and is worse off given the choice of A, then A comparatively harms her. If she has a bad life given A but does not exist given B, then A existentially harms

⁴ See, e.g., Bader 2022b; Cusbert & Kath 2018; McDermott 1982.

⁵ Thomas (2022) explicitly assumes a domain restriction along these lines; ? do not mention the topic of domain restriction, but we will charitably interpret them as including it. ? seems to be assuming a restriction of this kind in his defense of the claim that an option is impermissible only if it harms someone.

⁶ There are at least two other types of harm discussed in the literature on harm, but they are not relevant to our discussion. First, an option might impose non-comparative harm on an individual by creating her in an intrinsically bad state, or by creating her in a state which has an intrinsically bad aspect (Harman 2009; Shiffrin 1999). Second, an option might impose harm on an individual by making her worse off than she could have been in some specified possible outcome, where this outcome need not be one of the agent’s options. See for a discussion of different types of harm.

her.⁷ The harm done by a chosen option is defined as comparative or existential relative to some alternative that the agent could choose. It is therefore possible that an option A comparatively harms someone relative to some alternative B, but existentially harms the same person relative to some third alternative, C.

For instance, suppose I face the decision represented in Table 1. The lifetime well-being levels of people affected by my choice are represented numerically in the tables, where positive and negative numbers represent, respectively, positive and negative well-being, 0 represents neutral well-being, and ‘ Ω ’ represents non-existence.⁸

Table 1. Improvable Life

	Pebbles
Option 1	Ω
Option 2	-10
Option 3	10

In *Improvable Life*, Pebbles is existentially harmed by Option 2 relative to Option 1, since, given Option 2, Pebbles has a bad life, and she does not exist given Option 1. But Pebbles is comparatively harmed by Option 2 relative to Option 3, since, given Option 2, she is worse off than she is given Option 3. Following Jacob Ross, we will say that Pebbles has an *improvable life* in the outcome of Option 2. A person has an improvable life given the choice of some option, if the chosen option comparatively harms her.

How much harm should we say Option 2 does in *Improvable Life*? One possibility is that the harm of an option has a certain magnitude relative to some alternative, but there is no such thing as the magnitude of the option’s harm *full stop*. For instance, one might claim that Option 2 does existential harm of magnitude 10 relative to Option 1, and comparative harm of magnitude 20 relative to Option 3, but there are no further facts regarding how much harm Option 2 does. The theories considered in §5 are of this sort; they assume the magnitude of any morally significant harm is determined only relative to some alternative.⁹

Another possibility is that there is such a thing as the magnitude of an option’s harm *full stop*. For instance, we could say that the magnitude of any comparative harm *full stop* is the difference between the harmed individual’s well-being in the

⁷ On this distinction, see Bykvist (2006); McMahan (1981, 2013), Parfit (2017), Podgorski (2023), and Thomas (2022).

⁸ Since at least Boonin (2014), characters from Hanna-Barbera cartoons have sometimes featured in debates in population ethics. We continue the tradition here.

⁹ The qualification ‘morally significant’ is important. Technically, the theories considered in §5 leave open whether there is harm *full stop*; but if there is, then it is not morally significant on these theories.

outcome in which she is comparatively harmed and her well-being in the outcome in which, among the outcomes the agent can bring about, she is best off. On the other hand, if the only outcome in which the harmed person does not have negative well-being is an outcome in which she does not exist, then the magnitude of the existential harm imposed on her is simply her negative well-being level. Those theories considered in §§3–4 are of this sort.

On any harm-avoidance theory, the total harm of an option is the sum of all individual comparative and existential harms done by the choice of that option; this is either total harm full stop or total harm relative to some alternative, depending on one's theory.

Finally, we acknowledge that the conception of 'harming' employed in this paper is somewhat non-standard, and may strike some readers as odd. There is an understanding of 'harm', as a verb, according to which whether the choice of some option *harms* someone depends on whether the option involves "doing" as opposed to a merely "allowing". According to this understanding, if I choose an option that merely allows a certain person to drown, then I have allowed a morally significant harm, but I have not *harmed* the person, since I did not directly cause her drowning.¹⁰ Given our terminology, however, I *do* harm this person by allowing her to drown. Specifically, assuming she would have been better off had I done something other than allow her to drown, I comparatively harm her. This non-standard use of the active verb 'harm' serves mainly to *abbreviate* discussions of the cases we are interested in. Readers who find it odd that a mere allowing *harms* someone should feel free to interpret sentences such as 'I harmed the person by letting her drown' in some other way, such as 'By letting the person drown, I brought about an outcome in which she was harmed'. What matters for our purposes is that some options for an agent result in harm that certain alternatives for that same agent avoid. It doesn't matter *substantively* whether we describe these options as *harming*. Indeed, the moral relevance of the doing-allowing distinction is yet another deontic consideration that is typically bracketed in discussions of normative population ethics.

2.3 Conceptual framing

Harm-avoidance theories are often couched in terms of 'complaints', or 'objections' on behalf of a person who suffers harm as a result of the agent's choice.¹¹ One reason for this is related to our last point in the previous sub-section, i.e., the apparent oddness of describing mere allowings as harming. Rather than say that the choice of

¹⁰ Or so many assume, such as F. M. Kamm (1996); cf. Kagan 1991; Otsuka 1997.

¹¹ While harms may not be the only source of complaints, harms are a necessary source of complaints on the theories we discuss.

some option harms a person, when the option is a mere allowing, some prefer to say that the choice of that option gives the person who thereby suffers harm a complaint against the agent, or an objection to how the agent has behaved, where the basis of this complaint or objection is the harm that this person thereby suffers.

Another reason for adopting the language of 'complaints' or 'objections' is that, as remarks, this language pairs well with certain general moral theories, such as T. M. Scanlon's contractualism, as a matter of what people owe to each other.¹²

Couching a harm-avoidance theory in terms of complaints or objections also allows for an explanation of the Asymmetry that maps neatly onto the harm-avoidance account. Someone who is made to exist with a miserable life when the agent could have refrained from creating them has been existentially harmed, and therefore has a complaint. However, if the agent chooses not to create a person with a good life, then assuming no one else is harmed by the agent's choice, there is no one who can reasonably complain.

2.4 The moral significance of existential benefits

The final point of clarification concerns what a harm-avoidance theory takes to be the moral significance of creating people with good lives, i.e., the significance of conferring *existential benefits*. In a choice between options A and B, A existentially benefits someone if A causes her to exist with a good life but she would not exist given B. Like existential harm, the existential benefit of an option is defined relative to an alternative. The magnitude of an existential benefit to a person is simply her positive well-being level.

All harm-avoidance theories agree that agents have no moral obligation to create existential benefits rather than not create them, all else being equal. However, some harm-avoidance theories entail that creating existential benefits can justify harm, making certain otherwise impermissible harmful options permissible.¹³ For instance, it seems permissible to prolong the human race even though at least some future people will have bad lives, and hence, will suffer existential harm. Similarly, it seems permissible for parents to sacrifice some of their well-being for the sake of creating a happy child. One way to capture these intuitions is to claim that conferring large enough existential benefits can justify imposing harm.¹⁴

¹² The focus on complaints also helps bring out one implausibility of the *Problem of Improvable Life Avoidance*: what Horton (2021) refers to as 'backfiring complaints' (cf. §3, below).

¹³ Cf. McMahan, 2013; Thomas 2022; Podgorski 2023.

¹⁴ See Thomas (2022, §§4.2-4.3, as well as [removed for blind review], for discussion of these cases. See also Horton (2021, §1.2 for a discussion of how the claim that creating existential benefits can justify harm avoids what he calls 'the Problem of Tyrannical Complaints'.

Finally, although all harm-avoidance theories say there is no requirement to create existential benefits rather than not create them, as we mention above, a harm-avoidance theory can recognize a requirement to create some people with good lives rather than other people with good lives. See Table 2.

Table 2. Non-identity Case

	Elroy	Judy
Option 1	1	Ω
Option 2	Ω	10

We classify non-identity cases as those in which an agent has at least two options, A and B, where A would create someone who would not exist given B, and B would create someone who would not exist given A. These cases involve a choice between creating different *contingent* people, i.e., people whose existence depends on the agent's choice. In contrast, *addition cases* are those in which an agent can choose whether to create some contingent person (or people), but cannot choose between creating *different* contingent people.

In *Non-identity Case*, you can either existentially benefit one contingent person (Elroy) or existentially benefit a different contingent person (Judy) even more. It seems that creating Elroy (Option 1) does not harm Judy, and creating Judy (Option 2) does not harm Elroy.¹⁵ Since neither option does harm, many who defend what we classify as a harm-avoidance theory would accept the claim that either option is permissible. They would therefore reject the following intuitively plausible principle:

Normative Egalitarian Dominance (NED): For any options A and B, if the population that exists given the choice of A has perfect equality of welfare, is the same size as the population that exists given B, and every person who exists given A has higher welfare than every person who exists given B, then B is impermissible, other things being equal.¹⁶

However, our classification of harm-avoidance theories is broad enough to include theories that accommodate NED. In a choice between A and B, where A creates some contingent person S who would not exist given B, and B creates some contingent person S* who would not exist given A, let us say that A creates a *non-identity shortfall*

¹⁵ But see Meacham (2012) for an opposing view.

¹⁶ The title of this principle is originally due to Arrhenius (2022, p. 191). Our statement of the principle is similar to Arrhenius's, except that in ours, 'A' and 'B' denote possible options for an agent rather than populations, and we use the term 'impermissible' rather than 'wrong'.

if A existentially benefits S less than B existentially benefits S*. On our broad classification, a harm-avoidance theory can recognize both non-identity shortfall and harm imposition as sources of impermissibility. For instance, Thomas (2022) presents such a theory, which is compatible with NED (cf. §5.1). In a similar vein, Johann Frick and Michael Otsuka each propose a set of principles that reconciles the Asymmetry and NED.¹⁷

A harm-avoidance theory, then, is any theory that includes the following commitment:

Harmless Permission: If option A does no comparative or existential harm, and does not create any non-identity shortfall, then A is permissible.

Where the only considerations assumed to be relevant to permissibility are the number, identities, and well-being levels of people.

3. The Problem of Improvable Life Avoidance

Harmless Permission is compatible with a wide range of different general theories. The simplest harm-avoidance theory, which often serves as a starting point in discussions of how best to develop a harm-avoidance theory, is

Harm Minimization: An option A is permissible iff there is no alternative B which does less total harm than A.

The magnitude of a comparative harm is here assumed to be the difference between the harmed person's well-being in the outcome in which she is harmed and her well-being in the outcome in which, among the outcomes the agent could have brought about, she is best off. And the magnitude of an existential harm is assumed to be the harmed person's negative well-being level, where the only alternatives to existentially harming the person involve not creating her at all.

Harm Minimization implies that it is impermissible to create a miserable person rather than not create them, other things being equal. Creating the person would impose some existential harm, but refraining from creating the person would impose no harm, so not creating the person would do less harm than creating them. Harm Minimization also implies that it is permissible not to create a happy person rather

¹⁷ Frick (2020) defends the moral reasons formulation of the Asymmetry, not the deontic formulation. However, the moral reasons formulation supports the deontic formulation, and Frick presumably accepts the latter in addition to the former.

than create them, other things being equal. In this case, not creating the happy person would do no harm, so there can be no alternative that does even less harm.

However, as several philosophers have pointed out, Harm Minimization faces serious problems, one of which is *the Problem of Improvable Life Avoidance*.¹⁸ Table 3 illustrates a case, due to Jacob, which is commonly used to introduce the problem.¹⁹

Table 3. Ross's Case

	Roxy	Chip
Option 1	1	Ω
Option 2	10	10
Option 3	-2000	1000

In *Ross's Case*, Option 1 imposes a comparative harm of 9 on Roxy, Option 2 imposes a comparative harm of 990 on Chip, and Option 3 imposes a comparative harm of 2010 on Roxy. Option 1 minimizes total harm. Harm Minimization therefore implies that Options 2 and 3 are impermissible, and that Option 1 is morally required, as it is the only permissible option.

But this assignment of deontic statuses to options 1–3 seems implausible. Although Option 3 is clearly impermissible, the claim that Option 2 is also impermissible, and that Option 1 is therefore morally required, is problematic for at least two reasons.

First, it implies that the agent is morally required to avoid creating a certain person with a good life just because this life would be improvable. This is where *the Problem of Improvable Life Avoidance* gets its name. If Option 2 is impermissible, this can only be because it gives Chip an improvable life. Yet, Option 2 also gives Chip a *good* life. If Option 2 is impermissible only because it gives Chip a good but improvable life, it may seem odd that one is morally required not to create Chip. Presumably, Chip should prefer existence with a good life to non-existence. For instance, if we choose Option 2, and Chip objects that we have harmed him, we can respond “the only alternative for us that would not inflict even greater harm on someone else would leave you out of existence altogether. Is that really what you want?” One imagines that his answer would be “No”.²⁰

Notice also that according to Harm Minimization, in a binary choice between Options 1 and 2, Option 2 minimizes harm. So in this binary choice, Option 2 is morally required, and hence permissible, whereas Option 1 is impermissible, and hence, *not*

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

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

²⁰ This is an instance of what Horton dubs a ‘backfiring objection’ (cf. footnote 11). See also McDermott (2019).

morally required. But on Harm Minimization, Option 2 *becomes* impermissible, and Option 1 *becomes* morally required, when the horrible Option 3 is added to the option set. Let us say that a moral consideration against some option is a possible source of that option being impermissible. Then Harm Minimization violates

Improvable Life Acceptance (ILA): If (i) person S has a good life given A, (ii) does not exist given B, and (iii) the only moral consideration against A in a choice from some option set \mathcal{O} that includes A and B is that A comparatively harms S, then if B is not morally required in a binary choice between A and B, then B is not morally required in a choice from \mathcal{O} .

Basically, ILA says that just to avoid giving someone a good but improvable life, one is not morally required to leave that person out of existence.²¹

While ILA strikes us as fairly plausible, some would reject it on the grounds that in certain cases, refraining from creating a person with a good but improvable life is the only way to avoid unjust harm.²² Perhaps one reason why Harm Minimization's violation of ILA seems implausible when considering *Ross's Case* is that in this case the only option that gives Chip a better life than Option 2 is Option 3, which is unspeakably horrible for Roxy, and is, as Otsuka (2017) would say, "manifestly unreasonable". One might think that it is this detail, and not a principle such as ILA, that explains why the comparative harm that Option 2 imposes on Chip is insufficient to make Option 2 impermissible in a choice between Options 1–3.

Whatever one thinks about ILA, *Ross's Case* illustrates a second problem with Harm Minimization, namely that it morally requires dominated options. One option dominates another *iff* it is better for someone and worse for no one. In *Ross's Case*, Option 2 dominates Option 1, since it is better for Roxy and worse for no one. In fact, Option 2 *weakly addition-dominates* Option 1. Option A *addition-dominates* option B *iff* (i) everyone who exists given B would be better off if A were chosen, (ii) A creates

²¹ The following statement by Podgorski comes close to capturing the core idea of ILA: "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). There are only two differences between Podgorski's statement and our formulation of ILA; first, Podgorski refers to a "complaint" on behalf of the person mentioned, and second, he speaks of the option of creating the person being 'permitted' or 'impermissible', whereas we speak of the option of *not* creating the person as being 'not morally required'. Our formulation is weaker than Podgorski's, since it doesn't assume anything about complaints, and it leaves open the (admittedly implausible) possibility of the agent facing a moral dilemma when C is added to the option set alongside A and B.

²² See, e.g., Boonin (1996) and Frick (2022), as well as Temkin (2012, ch. 13). Boonin and Frick discuss Parfit's *Mere Addition Paradox* as a case where adding people with good but improvable lives results in a morally worse outcome relative to the option set. Ingmar Persson (2017, ch. 8) argues that it can be worse to add people with good but improvable lives if this increases inequality; and Temkin (2012, ch. 12) suggests that this may be the case.

some people who would not exist given B, and (iii) everyone who exists given A has a good life. On the other hand, A *weakly* addition-dominates B *iff* A addition-dominates B and everyone who exists given A has equal well-being. Since Harm Minimization implies that Option 1 is morally required in *Ross's Case*, it violates

Weak Dominance Addition Exemption (WDAE): If option A weakly addition-dominates option B, then B is not morally required.²³

Most harm-avoidance theorists seem to agree we should accept WDAE.²⁴

We have two problems here for those who wish to develop a general harm-avoidance theory. First, the simplest harm-avoidance theory violates ILA (the Problem of Improvable Life Avoidance); second, it violates WDAE. For harm-avoidance theorists, there are different ways of responding to these problems. They could find a compelling justification for rejecting ILA and WDAE, and perhaps supply an alternative explanation of where Harm Minimization goes wrong in *Ross's Case*, such as that comparative harm doesn't count against an option when the alternative that is better for the harmed person is "manifestly unreasonable". They could formulate a harm-avoidance theory that satisfies both ILA and WDAE without violating some equally compelling principle. Or they could adopt a mixed approach that involves formulating a theory that accommodates only one of the two principles, while supplying a justification for rejecting the other. McDermott (2019), Horton (2021), and Podgorski (2023) have taken the second approach, while Thomas (2022) adopts a combination of the second and third approaches, offering one theory that satisfies WDAE and ILA, and a second theory that satisfies WDAE but not ILA.

However, as we will now argue, these harm-avoidance theories have other problematic implications, some even more implausible than violating WDAE, and others even more implausible than violating ILA.

²³ Our formulation of this principle is inspired by Elliot Thornley, who appeals to a weaker principle in arguing against a theory defended by Joe Horton, which we consider in §3. Thornley (2023, p. 522) calls his principle 'Weak Normative Dominance Addition', which is like WDAE, except it says that if everyone has non-negative well-being in the weakly addition-dominated option, then if that option is permissible, the weakly addition-dominating option is also permissible.

²⁴ Some seem to believe that dominated options can be required. For example, Frick (2022, 238ff) suggests that in one "supercharged" version of the *Mere Addition Paradox*, where one of the options is dominated, the dominated option might be morally required. Frick explicitly argues only for the claim that the dominated option is better than the dominating option relative to a certain set of options; but his discussion suggests that the dominated option is also the best option in this set, and that given the absence of any non-axiological considerations, it is required.

4. Avoid Reasonable Objections

McDermott (2019) and Horton (2021) each defend a harm-avoidance theory that accommodates both ILA and WDAE. Although McDermott's and Horton's 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 be harmed by it.

On McDermott's theory, which he calls '*Objection Minimization*', an individual S has a reasonable objection to an option A *iff* S exists given A, and there is some alternative B such that (i) B is better for S than A (and hence, A harms S), and (ii) B does less total harm than A. Objection Minimization implies that in *Ross's Case* Chip does not have a reasonable objection to Option 2, since the only option that is better for Chip, Option 3, does more total harm than Option 2. It also implies that Roxy does not have a reasonable objection to Option 2, since there is no alternative that is better for her. So according to Objection Minimization, Option 2 is permissible.

One drawback of Objection Minimization is that it implies that Option 1 is permissible.²⁵ Since Option 2 does more total harm than Option 1, Objection Minimization implies that Roxy cannot reasonably object to Option 1. Since Roxy is the only potential objector to Option 1, no one can reasonably object to Option 1. So Objection Minimization violates

Weak Dominance Addition (WDA): If A weakly addition-dominates option B, then B is impermissible.

However, permitting weakly addition-dominated options is not as implausible as requiring them. So Objection Minimization seems at least to improve upon Harm Minimization.²⁶

Horton's harm-avoidance theory implies that in *Ross's Case* Option 2 is morally required and Options 1 and 3 are impermissible, which seems correct. Horton's formulation of his criteria for reasonable objectionableness are somewhat complicated. For ease of exposition, our statement of the criteria differs slightly from Horton's, but this doesn't affect our arguments.²⁷

According to Horton's *Avoid Reasonable Objections*:

²⁵Thomas (2022, fn 8) makes this point.

²⁶ Cf. Thomas 2022, fn 23.

²⁷ His original statement can be found at (Horton 2021, p. 499).

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—4.²⁸

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 given A.
3. The sum of well-being that B gives to the set of people who exist given B is greater than the sum of well-being that A gives to the set of people who exist given A.²⁹
4. No one can reasonably object to B.

An option is permissible *iff* no one can reasonably object to it.

Two clarifications are needed. First, Horton thinks that B can be better or worse for someone than A, even if she does not exist given B. If she has a bad life given A, but does not exist given B, then on Horton's view, assuming she exists (i.e., A has been chosen), B is better for her than A. Second, for the purpose of determining whether condition 2 is satisfied, the sum of well-being that B gives to the set of people who exist given A is the sum of the individual well-being values *of B* for those who exist given A, where some people who exist given A may not exist given B. If someone who exists given A does not exist given B, then, Horton assumes, B gives *zero* well-being to this person.

According to Avoid Reasonable Objections, in *Ross's Case*, Option 2 is morally required because it is the only option no one can reasonably object to. To see this, notice that neither Roxy nor Chip can reasonably object to Option 2. Chip cannot reasonably object to Option 2 because his objection cannot satisfy conditions 2 and 3. The only option that is better for Chip than Option 2 is Option 3, which produces less well-being than Option 2 for the set {Roxy, Chip}. Roxy cannot reasonably object to Option 2 because her objection cannot satisfy condition 1, i.e., there is no alternative to Option 2 that is better for Roxy.

²⁸ Although it may be unclear from the four conditions stated here, Avoid Reasonable Objections is indeed a harm-avoidance theory. According to this theory, the permissibility of an option requires that no one could reasonably object to it, and a necessary condition for reasonably objecting to an option is that there is some alternative that is better for the objector. Where A is the option to which the objector objects, if the relevant alternative B is better for the objector than A because the objector has a bad life given A and does not exist given B, then the objector is existentially harmed by A; on the other hand, if B is better than A for the objector because she would exist at a higher level of well-being given B, then the objector is comparatively harmed by A. So on Avoid Reasonable Objections, an individual has a reasonable objection to A only if A harms her.

²⁹ We've included Thornley's (2023, p. 519) amendment to Horton's condition 3 in our statement.

On the other hand, Roxy can reasonably object to Option 3. Option 2 is better for Roxy than Option 3, so condition 1 is satisfied. Option 2 also produces a greater sum of well-being than Option 3 for the set {Roxy, Chip}. This is sufficient for Roxy's objection to satisfy conditions 2 and 3 because Roxy and Chip are the only people who exist given either Option 2 or Option 3. Finally, because no one can reasonably object to Option 2, Roxy's objection to Option 3 satisfies condition 4.

Roxy can also reasonably object to Option 1, since Option 2 is better for her, gives a greater sum of well-being than Option 1 to the set {Roxy}, the sum of well-being that Option 2 gives to the set {Roxy, Chip} is greater than the sum of well-being that Option 1 gives to the set {Roxy}, and no one can reasonably object to Option 2.

Since someone can reasonably object to Options 1 and 3, and no one can reasonably object to Option 2, Option 2 is the only permissible option, according to Avoid Reasonable Objections.

Although Avoid Reasonable Objections provides a plausible treatment of *Ross's Case*, and does not require weakly addition-dominated options, Horton acknowledges that it sometimes *permits* weakly addition-dominated options, thereby violating WDA. He also acknowledges that Avoid Reasonable Objections sometimes implies that when one option weakly-addition dominates another, the latter is permissible and the former impermissible, an objection emphasized by Thornley (2023).

Horton illustrates this with the following case, which he thinks demonstrates a *strength* of Avoid Reasonable Objections:³⁰

Table 4. Horton's Case 6

	Barney	Betty
Option 1	1	Ω
Option 2	2	2
Option 3	Ω	100

In *Horton's Case 6*, Option 2 weakly addition-dominates Option 1. But Avoid Reasonable Objections implies that Option 2 is impermissible and that Option 1 is permissible. Option 2 is impermissible because Betty has a reasonable objection to it. She exists given Option 2, she does not (Horton assumes) consent to this act, and there is an alternative to Option 2, namely Option 3, which is better for Betty, produces more well-being for the set {Betty, Barney}, and produces more well-being for {Betty} than Option 2 produces for {Betty, Barney}. Moreover, no one has a reasonable objection to Option 3 according to Avoid Reasonable Objections, since Betty is the only person who exists given Option 3, and there is no alternative to Option 3 that is better for

³⁰ Our presentation of his case was sourced, with minor cosmetic changes, from Horton (2021, p. 496).

her. So Betty's objection to Option 2 satisfies all four of Horton's conditions for reasonableness. Because Betty has a reasonable objection to Option 2, Barney *does not* have a reasonable objection to Option 1. The only option that would be better than Option 1 for Barney, namely Option 2, is one that Betty can reasonably object to; so Barney's objection to Option 1 does not satisfy condition 4. Since no one has a reasonable objection to Option 1, according to Avoid Reasonable Objections, Option 1 is permissible in *Horton's Case 6*.

The fact that Avoid Reasonable Objections implies Option 1 is permissible but Option 2 impermissible seems like a problem, though again, this problem is not as grave as that of requiring weakly addition-dominated options.³¹ Intuitively, Option 3 is *morally required* in *Horton's Case 6*. Not only does Option 3 produce the most well-being of any option, but more importantly, it is the only option that avoids harm. However, Horton thinks we should reject the claim that Option 3 is morally required because, on Avoid Reasonable Objections, this claim violates ILA. Recall that according to ILA, if, in a choice between creating someone with a good life (A) and leaving them out of existence (B), we are not required to choose B, then adding another option (C) that is better for this person than A cannot generate a moral requirement to choose B. To see why requiring Option 3 would violate this principle on Avoid Reasonable Objections, consider the following. According to Avoid Reasonable Objections, in *Horton's Case 6*, Option 3 would not be morally required in a binary choice between only Options 1 and 3, since, in such a binary choice, neither option would harm anyone. If Option 3 is not required in a binary choice between Options 1 and 3, but Option 3 *becomes* required when Option 2 is added to the option set, then we have a straightforward violation of ILA. Barney exists with a good life given Option 1, he does not exist given Option 3, and the only moral consideration against Option 1 is that it comparatively harms Barney, since Option 2 is better for him than Option 1.

Two points are worth emphasizing here. First, intuitively, Option 3 is morally required in a binary choice between Options 1 and 3. In this binary choice, Option 1 creates one person (Barney) rather than a different person (Betty) with a much better

³¹ Note that Avoid Reasonable Objections violates Thornley's Weak Normative Dominance Addition principle. It is worth emphasizing that in *Horton's Case 6*, the possible well-being levels for those who might exist given the different options seem to be chosen to minimize the intuitive implausibility of Avoid Reasonable Objections's violation of WDA. These levels are chosen so that the difference in well-being for Betty between Options 2 and 3 is relatively large, while the difference for Barney between Options 1 and 2 is relatively small. But as Thornley (2023, p. 522) points out, Avoid Reasonable Objections would assign the same deontic statuses to the three Options in any case with a similar structure but where the well-being difference for Barney between Options 1 and 2 is much greater, and just barely smaller than the well-being difference for Betty between Options 2 and 3. (For instance, imagine, following Thornley's example, that both Barney and Betty would have well-being 49 in the outcome of Option 2.)

life. Since Avoid Reasonable Objections denies this, it violates NED.³² More generally, it fails to account for the apparent moral significance of non-identity shortfall. Horton would not be moved by this objection, since he bites the bullet in response to *the Non-identity Problem*. However, a harm-avoidance theorist could try to modify Avoid Reasonable Objections to account for the moral significance of non-identity shortfall. For instance, one could keep Horton's criteria for an objection being reasonable, but modify his criterion of permissibility to allow for the possibility of an option being impermissible when it causes non-identity shortfall. Of course, one would then need to figure out how to fit these different criteria together into a coherent moral theory. The theoretical benefit would be that one could say that Option 3 is morally required in *Horton's Case 6* without rejecting ILA.

Second, even if a harm-avoidance theorist insists on rejecting NED, she could reasonably view *Horton's Case 6* as a counterexample to ILA. Even if one claims that Options 1 and 3 are both permissible in a binary choice, one could justify the claim that Option 1 becomes impermissible when Option 2 is added to the option set on the grounds that Option 1 then harms someone. Specifically, one could claim that if an agent must create someone with a good life, and the agent's choice is between creating a person with a good but improvable life and creating a different person with a life that is at least as good but *unimprovable*, then the agent ought to create the latter. Violating ILA in this type of case may not seem too high a cost for securing the plausible judgment that Option 3 is morally required in *Horton's Case 6*.

Unfortunately, regardless of how one addresses the foregoing points, Avoid Reasonable Objections faces a further serious objection. This objection applies to McDermott's Objection Minimization as well. Both theories violate

Weaker Dominance Addition (Weaker DA): If A weakly addition-dominates B, and everyone who exists given B has a bad life, then B is impermissible.

Permitting weakly addition-dominated options that give everyone a bad life seems even more implausible than requiring weakly addition-dominated options which, like Option 1 in *Ross's Case*, at least give everyone a good life.

To see that Avoid Reasonable Objections and Objection Minimization violate Weaker DA, consider the following case:

³² McDermott's Objection Minimization also violates NED, since his theory implies that someone can reasonably object to an option only if it harms her, whereas NED implies that an option is impermissible if it creates a person with less well-being than some other person whom one could instead have created, even if this option does not harm anyone.

Table 5. Weaker Dominance Addition Violation

	Wilma	Fred
Option 1	-100	Ω
Option 2	100	100
Option 3	-100	1000

In this case, everyone who exists given Option 1 has a bad life, everyone who exists given Option 2 has a good life, and Option 2 weakly addition-dominates Option 1. Yet both Avoid Reasonable Objections and Objection Minimization imply that Option 1 is *permissible*.

According to Objection Minimization, Option 2 does more total harm than Option 1; Option 2 imposes a comparative harm of 900 on Fred, while Option 1 imposes a comparative harm of only 200 on Wilma. Moreover, there is no alternative to Option 1 that is both better for Wilma and does less total harm than Option 1. Option 3 imposes a comparative harm of 200 on Wilma, and hence does the same amount of harm as Option 1; moreover Option 3 is not better than Option 1 for Wilma. So according to Objection Minimization, Wilma’s objection to Option 1 is not reasonable. Hence, according to Objection Minimization, Option 1 is permissible.

To see how the same problem arises for Horton’s theory, notice that according to Avoid Reasonable Objections, no one has a reasonable objection to Option 3. If anyone had a reasonable objection to Option 3, it would be Wilma. But the only alternative that is better than Option 3 for Wilma is Option 2, which produces less well-being than Option 3 for {Wilma, Fred}. So Wilma’s objection to Option 3 cannot satisfy conditions 2 and 3. Given that no one has a reasonable objection to Option 3, Fred has a reasonable objection to Option 2; he exists given Option 2, does not (we assume) consent to Option 2, and there is an alternative, Option 3, which is better than Option 2 for Fred, produces more well-being than Option 2 for the set {Wilma, Fred} and produces more well-being for {Wilma, Fred} than Option 2 produces for {Wilma, Fred}. Finally, since Fred has a reasonable objection to Option 2, it follows that Wilma’s objection to Option 1 cannot satisfy condition 4. She has no reasonable objection to Option 1, making Option 1 permissible according to Avoid Reasonable Objections.

This problem cannot be avoided by modifying Avoid Reasonable Objections and Objection Minimization to accommodate the apparent moral significance of non-identity shortfall. This is because *Weaker Dominance Addition Violation* is not a non-identity case, but an addition case. There is no pair of options where one involves creating a certain contingent person and the other involves creating a different contingent person.

Horton and McDermott might try to avoid the implausible implication that Option 1 is permissible by adopting a prioritarian weighting that assigns greater moral weight to harm-based objections the worse off the objector is.³³ This would allow one to say that Wilma's objection to Option 1 carries more weight than Fred's objection to Option 2. However, this move won't help, since we can imagine the well-being difference for Fred between Options 1 and 2 to be as great as we want. Unless some harm-based objections are infinitely more morally weighty than others, where both objections are based on finite amounts of harm, there will be some magnitude of harm that we can imagine Option 2 imposing on Fred such that his objection to Option 2 will outweigh Wilma's objection to Option 1. We can also imagine a case with the same structure as *Weak Dominance Addition Violation* in which *any number of people* in Fred's position would be harmed by Option 2 to the same extent as Fred. It is hard to see how Wilma's harm-based objection to Option 1 can outweigh any number of Fred-like objections to Option 2.

A different response on behalf of Horton's and McDermott's theories would be to modify these theories to allow for the possibility that someone can have a reasonable objection to an option even when the only alternative that is better for them does more total harm, or is such that someone else can reasonably object to it. This would allow that Wilma's objection to Option 1 is reasonable even though Fred has a reasonable objection to Wilma's preferred alternative, Option 2. However, modifying either theory in this way would give up on its treatment of *Ross's Case*. Specifically, it would open the door to the possibility that in *Ross's Case*, each option is one that someone can reasonably object to. This would make *Ross's Case* a moral dilemma, which seems absurd.

Finally, we note that Harm Minimization also implies that Option 1 is permissible in *Weaker Dominance Addition Violation*, since no alternative to Option 1 does less harm than Option 1. Thus, all three theories we've considered so far violate Weaker DA, which is, we think, worse than violating WDAE.

5. Tournament Theories

Unlike Harm Minimization, Objection Minimization, and Avoid Reasonable Objections, the two harm-avoidance theories considered in this section satisfy WDAE, WDA, and Weaker DA. They forbid weakly addition-dominated options.

Both Thomas's and Podgorski's theories adopt what Podgorski calls a *tournament approach*, and we shall refer to them as 'tournament theories'. A tournament theory is

³³ Horton (2021, n. 6, 17) considers this kind of view.

structured in two parts. First, it includes a set of conditions for when one option defeats another option, or when one ought to choose the former over the latter, in a binary choice, or pairwise comparison, of only those two options. Second, it includes a further condition that determines when the choice of some option from any finite set of options is (im)permissible, where this determination is based on how each option fares in pairwise comparisons with the other options.

In contrast to the theories considered in §§3–4, on the tournament approach, what matters is how much harm an option does in a pairwise comparison with each alternative (cf. §2.2). For instance, if our options are A, B, and C, we must ask how much harm A does in a pairwise comparison with B (ignoring the presence of C), then how much harm A does in a pairwise comparison with C (ignoring the presence of B), and then apply the same procedure to Options B and C. Depending on how each option fares in a pairwise comparison with the others, our general criterion of permissibility will then give us the deontic status of each option.

To illustrate, recall *Ross's Case*, in which, according to Harm Minimization, Option 2 does more harm than Option 1, despite the fact that Option 2 weakly addition-dominates Option 1. In contrast with Harm Minimization, a tournament approach says that the harm done by Option 2 is relevant only in a pairwise comparison of Options 2 and 3. Since Option 3 is clearly horrible, any sensible tournament theory will imply that Option 2 defeats Option 3, or that Option 2 ought to be chosen over Option 3, in a binary choice between Options 2 and 3. But notice also that in a binary choice between Options 1 and 2, Option 1 does more harm than Option 2, since, in that binary choice, Option 1 harms Roxy but Option 2 harms no one. Option 2 therefore “wins” in a pairwise comparison with either Option 1 or Option 3. Thus, any sensible tournament theory will imply that Option 2 is at least permissible. In fact, on Thomas's and Podgorski's theories, Option 2 is morally required in *Ross's Case*.

However, as we will now argue, each of these theories encounters problems which stem, at least in part, from the tournament approach.

5.1 The Maximization Theory

The first tournament theory we consider is due to Thomas (2022).

Thomas actually presents two theories that differ regarding their treatment of *the Non-identity Problem*. The first, which Thomas calls ‘a narrow theory’ rejects NED, biting the bullet in response to *the Non-identity Problem*. The second, which he calls ‘a

wide theory', accommodates NED. He considers *the Non-identity Problem* so vexed that he leaves it an open question which of these theories is more plausible.³⁴

Each theory includes its own criteria for when one *ought to choose* some option A over another B in a binary choice between A and B. Because the objection that we will raise applies to both the narrow and wide theories, we will here consider only the narrow theory, which is the simpler of the two theories.

The conditions governing pairwise comparisons on the narrow theory are presented as follows. In a binary choice between any options A and B, let Harm(A) be the total harm (both comparative and existential) that arises from choosing A over B. Let ExBen(A) be the total existential benefit in A, and ExBen(B) the total existential benefit in B. Then one ought to choose A over B *iff*:

1. Harm(B) > Harm(A)
2. Harm(B) + ExBen(A) > Harm(A) + ExBen(B)

The motivation for conditions 1 and 2 is as follows. First, we have moral reasons to avoid comparative and existential harm. These reasons have requiring strength proportionate to the magnitude of the harm that would be inflicted. We also have moral reasons to create existential benefits. However, these reasons have *no* requiring strength; they have only justifying strength.³⁵ They can defuse competing requiring reasons to avoid harm, but they cannot by themselves generate moral requirements. The justifying strength of one's reason to existentially benefit someone is proportionate to the magnitude of the existential benefit.³⁶ The narrow theory's condition 1 reflects the idea that there is requiring moral reason to avoid harm, and that in a binary choice between two options it is never the case that the agent ought to choose the option that she has more requiring reason not to choose, i.e., the option that does greater harm. Condition 2 reflects the idea that the purely justifying moral reason to existentially benefit people can neutralize the requiring strength of the moral reason to avoid harm, but also that this purely justifying moral reason cannot by itself make it the case that an agent ought to choose one option over another in a binary choice. Notice, for example, that on the narrow theory, it is not the case that one ought to create a person with a very good life rather than some other person with a life that is good, but not very good. In this binary choice, neither option does less harm than the

³⁴ But see [removed for blind review] for arguments in favour of the narrow theory over the wide theory.

³⁵ Rebellious against the old fashion that reasons exclusively issue *pro tanto* requirements, philosophers are increasingly adopting the position that reasons can vary on at least two dimensions with respect to their normative strength (Gert 2004; Kaczmarek & Lloyd forthcoming; Kamm 1985; Lazar 2013; Munoz 2021; Pummer 2023). See esp. Little and Macnamara (2021) for an overview of this literature.

³⁶ Notice Thomas (2022, p. 490) crafted condition 2 to express the plausible idea that non-requiring reasons justify harm only when the *net* existential benefits favour that outcome.

other (since neither does any harm) and so neither satisfies the narrow theory's condition 1.

Thomas's criterion of permissibility for the narrow (as well as the wide) theory is:

Maximization: In a choice between finitely many options, all and only the maximal options are permissible.

That an option is 'maximal' means that *it is not the case* that one ought to choose some other option over it in a binary choice. Hence, whether A is permissible in a choice between finitely many options, depends on whether there is some B in the option set such that one ought to choose B over A in a binary choice. If so, then A is impermissible. Otherwise, A is permissible.

The narrow theory is the conjunction of Maximization and conditions 1 and 2. To see that the narrow theory satisfies WDA, consider any choice context in which A and B are options. If A weakly addition-dominates B, then the only difference between A and B, when compared pairwise, is that everyone who exists given B also exists given A with higher (positive) welfare, and A creates some additional people, also with positive welfare, who do not exist given B, such that everyone who exists given A is equally well-off. It follows that in a binary choice between A and B, the choice of B would impose some (comparative) harm, but would not create any existential benefits, while the choice of A would impose no harm and would create some existential benefits. Hence, when A weakly addition-dominates B, $\text{Harm}(B) > \text{Harm}(A)$, and $\text{Harm}(B) + \text{ExBen}(A) > \text{Harm}(A) + \text{ExBen}(B)$. So according to the narrow theory's conditions 1 and 2, one ought to choose A over B. Since one ought to choose A over B, it follows from Maximization that in any choice context that includes A and B, the choice of B is impermissible. WDA is satisfied, and since WDA entails WDAE and Weaker DA, the latter are also satisfied.

However, as we illustrate below, the narrow theory leads to a troubling form of improvable life avoidance that we call *strong improvable life avoidance*. Although we have not here considered the wide theory, the contexts in which the narrow theory leads to strong improvable life avoidance are those in which the narrow and wide theories agree on which options are (im)permissible. So strong improvable life avoidance is a problem for both theories. Using the label '*the Maximization Theory*' for the disjunction of the wide and narrow theories, our objection to the Maximization Theory is that it entails strong improvable life avoidance.³⁷

To illustrate, consider the following case:

³⁷ 'The Maximization Theory' is our label, not Thomas's.

Table 6. Strong Improvable Life Avoidance

	George	Jane
Option 1	100	Ω
Option 2	101	0
Option 3	-100	202

On the Maximization Theory, the deontic statuses of the options are those given by the narrow theory. They are determined as follows. First, consider a binary choice between Options 1 and 2. In this binary choice, Option 1 does more harm than Option 2. Moreover, since Jane is the only contingent person, and she has welfare 0 (a neutral life) given Option 2, neither option creates any existential benefits. Hence, one ought to choose Option 2 over Option 1.

Next, consider a binary choice between Options 2 and 3. In this binary choice, Option 2 does more harm than Option 3. Specifically, Option 2 imposes harm of 202 on Jane, while Option 3 imposes harm of only 201 on George. Moreover, neither option produces existential benefits, since the same people exist given either option. Hence, one ought to choose Option 3 over Option 2.

Finally, consider a binary choice between Options 1 and 3. Here, according to the Maximization Theory, neither option is such that it ought to be chosen over the other. Although Option 1 does less total harm than Option 3, it does not produce any existential benefits, while Option 3 produces an existential benefit for Jane that is larger than the comparative harm that Option 3 imposes on George. Thus, Option 3 is the only maximal option in this case, i.e., the only option such that no other option ought to be chosen over it in a binary choice. So by Maximization, Option 3 is the only permissible option, and is therefore morally required.

Notice that because neither Option 1 nor Option 3 ought to be chosen over the other in a binary choice, by Maximization, in such a binary choice, Option 1 is permissible and therefore Option 3 is *not* morally required. It is only when we add Option 2, which is (slightly) better for George than Option 1, to the option set, that the Maximization Theory requires Option 3, which is (much) worse for George than either Option 1 or Option 2. The Maximization Theory therefore violates

Weak Improvable Life Acceptance (WILA): If (i) A imposes greater harm on person S than B, and (ii) the only moral consideration against B, in a choice from an option set \mathcal{O} that includes A and B, is that B harms S, then if A is not morally required in a binary choice between A and B, then A is not morally required in a choice from \mathcal{O} .

As in our statement of ILA, here, by ‘the only moral consideration against B’ we mean the only potential source of B’s being impermissible.

The intuitive idea behind WILA can be grasped by first recalling the intuitive idea behind ILA: to avoid giving someone a good but improvable life, one is not morally required to leave her out of existence. In contrast, the idea behind WILA is that to avoid giving someone an improvable life, one is not morally required to instead give her a life that is *even more* improvable.

In the *Strong Improvable Life Avoidance* case, the only morally relevant consideration against Option 1, on the Maximization Theory, is that it comparatively harms George. But how can *that* generate a moral requirement to impose even greater harm on George by choosing Option 3? A requirement to choose Option 3 would be understandable if we had a requiring reason to existentially benefit Jane rather than leave her out of existence. But on the Maximization Theory, as on all harm-avoidance theories, there is no such requiring reason.

To see that WILA is more plausible than ILA, recall that, as we suggested in §2 and §3, someone could reject ILA on the grounds that in certain cases, in order to avoid comparatively harming someone, we can be required not to create them even with a good life. For instance, this might be the only way to avoid unjust harm. But this rationale for rejecting ILA does not support rejecting WILA. It is patently absurd to claim that just to avoid comparatively harming someone, we can be required to comparatively harm that same person *even more*.

The Maximization Theory violates WILA because it rules out options as impermissible on the basis of pairwise comparisons. For instance, Option 1 is deemed impermissible solely on the basis of a pairwise comparison with Option 2, and Option 2 is deemed impermissible solely on the basis of a pairwise comparison with Option 3. The Maximization Theory therefore cannot account for any potentially morally significant relations between Options 1—3 when all three options are considered together, for instance, the fact that the person who would be harmed by Option 1 in relation to Option 2 (George) is the same person who would be harmed even more by Option 3 in relation to either Option 1 or Option 2.

5.2 Minimize Unanswered Complaints

Thus far, we have seen that Harm Minimization, Objection Minimization, and Avoid Reasonable Objections satisfy WDAE and ILA but violate Weaker DA, and that the Maximization Theory satisfies all the dominance principles but violates WILA. This motivates the search for a harm-avoidance theory that accommodates both WILA and the dominance principles.

The second tournament theory that we shall consider, due to Podgorski (2023),

accomplishes this. According to his *Minimize Unanswered Complaints*, when an option causes existential or comparative harm to an individual, this provides grounds for a complaint on behalf of the individual against the choice of the option, where this complaint is had *relative* to some alternative that would have either made the individual better off or not harmed them.³⁸

Moreover, on this theory, existential benefits to individuals provide what Podgorski calls '*answers*' to complaints. When comparing only two options, A and B, the strength of an individual's complaint against A relative to B is the magnitude of the existential or comparative harm she incurs in the outcome of A. And if an individual exists conditional on A but not on B, then she generates an answer to harm-based complaints resulting from the choice of A *iff* her well-being conditional on A is positive, and the strength of this answer is the magnitude of her positive well-being.

When comparing only two options, A and B, if A harms some individual, then her complaint against A relative to B is unanswered, either entirely or partially, if the total of existential benefits brought about by A relative to B is less than the harm to this individual. And if there are unanswered complaints against A, then the total strength of the unanswered complaints against A relative to B is equal to the total harm of A relative to B minus the existential benefits of A relative to B.

Podgorski's theory includes both a criterion for when one option 'defeats' another in a pairwise comparison, and a general condition of permissibility, based on the criterion of defeat. He states his criterion of defeat as follows:

Minimize Aggregate Unanswered Complaints*: An option X defeats option Y *iff* the strength of unanswered complaints against X relative to Y is less than the strength of unanswered complaints against Y relative to X.³⁹

The general criterion of permissibility based on the above criterion of defeat is what Podgorski calls

Uncovered: An option is permissible *iff* there is no option that covers it, where A *covers* B *iff* A defeats B and any option(s) that B defeats.

³⁸ 'Minimize Unanswered Complaints' is our label, not Podgorski's.

³⁹ Notice that because a harm-based complaint is just as strong as the magnitude of the harm imposed on the complainant, and because the strength of an existential benefit answer is just as strong as the magnitude of that benefit, Podgorski's criterion of defeat can also be stated more simply in terms of harm and existential benefit, using Thomas's formalism:

Minimize Aggregate Unanswered Harm (MAUH): In a binary choice, A *defeats* B *iff* both (i) $\text{Harm}(B) - \text{ExBen}(B) > 0$ and (ii) $\text{Harm}(B) - \text{ExBen}(B) > \text{Harm}(A) - \text{ExBen}(A)$.

A *defeats* B just in case B has at least some unanswered harm and the total unanswered harm of B is greater than that of A.

Minimize Unanswered Complaints is the conjunction of Minimize Aggregate Unanswered Complaints* and Uncovered.

Like The Maximization Theory, Podgorski’s theory satisfies both WDA and Weaker DA. According to Uncovered, in any choice situation, B is permissible *iff* there is no option that covers B. But given Podgorski’s criterion of defeat, one can prove that in any choice context where A and B are both options, if A weakly addition-dominates B, then A covers B, i.e., for any C, if B defeats C, A defeats C. We prove this in the appendix.⁴⁰

Minimize Unanswered Complaints also satisfies ILA, as well as WILA. This is because, as Podgorski points out, an important property of Uncovered is that “losers cannot dislodge winners”; if some option A is permissible, then the addition of option B can make A impermissible only if B is permissible. A theory violates ILA and WILA only when it implies that the introduction of an impermissible option can flip the deontic status of one of the other options from permissible to impermissible. But if this cannot happen, then ILA and WILA are guaranteed.

Since Minimize Unanswered Complaints satisfies WDA, Weaker DA, and WILA, it may seem like a promising harm-avoidance theory.

However, Thornley (2023) has recently raised a serious objection to Minimize Unanswered Complaints. See Table 7.⁴¹

Table 7. Thornley’s Case

	Huckleberry	Yogi
Option 1	100	Ω
Option 2	0	2
Option 3	Ω	1

In a binary choice between Options 1 and 2, it is clear that Huckleberry has the strongest unanswered complaint against Option 2, and that therefore Option 1 defeats Option 2 according to Minimize Unanswered Complaints. However, simply introducing the possibility of making Yogi’s life worse (Option 3) makes Option 2 permissible, as now there is no option that covers Option 2. In other words, there is no option that defeats Option 2 *and* any option that Option 2 defeats. For according to Minimize Unanswered Complaints, although Option 1 defeats Option 2, it does not defeat Option 3, since no one is harmed by Option 3 relative to Option 1, or by Option 1 relative to Option 3. Even worse, as Thornley (2023) notes, Option 2 will be permissible no matter how strong Huckleberry’s harm-based complaint is against

⁴⁰ In the proof, the harm-based formulation MAUH is used (cf. footnote 41).

⁴¹ Our presentation of his case was sourced, with minor cosmetic changes, from Thornley (2023, p. 524).

Option 2 relative to Option 1, and no matter how little harm Option 2 prevents from befalling Yogi relative to Option 3. Thornley calls this ‘*the Problem of Impairable Life Acceptance*’.

Thornley’s Case demonstrates that Minimize Unanswered Complaints violates NED, a result which Podgorski, like Horton, is happy to accept. Option 3 creates a non-identity shortfall relative to Option 1, since Option 1 gives Huckleberry a life that is much better than the life that Option 3 gives Yogi; yet Minimize Unanswered Complaints entails that neither option defeats the other. One might therefore wonder whether *the Problem of Impairable Life Acceptance* could be avoided by modifying Podgorski’s criterion of defeat to reflect the apparent moral significance of non-identity shortfall as well as that of harm. Such a modified criterion would of course need to be worked out, and one would need to decide how non-identity shortfall is to be weighed against unanswered harm for the purpose of determining defeat. But the criterion would at least generate plausible results in *Thornley’s Case*. It would imply that Option 1 defeats Option 3 because of Option 3’s non-identity shortfall relative to Option 1, that Option 1 defeats Option 2 because of Huckleberry’s harm-based complaint against Option 2 relative to Option 1, and that Option 2 defeats Option 3 because of Yogi’s harm-based complaint Against Option 3 relative to Option 1. Option 1 would then cover both Option 2 and Option 3, and so Options 2 and 3 would be impermissible and Option 1 morally required, which is intuitively the correct result.

However, even such a revamped version of Minimize Unanswered Complaints would have the troubling feature that an option against which some person has an unanswered complaint, no matter how strong, can be permitted. Let us say that an amount of harm is unanswered *iff*, corresponding to that harm, there are unanswered complaints of a certain strength. Then Minimize Unanswered Complaints will sometimes permit any amount of unanswered harm, even if it is modified to account for the significance of non-identity shortfall.

For instance, consider Table 8, which represents a range of different possible addition (as opposed to non-identity) cases where x and y represent different well-being values that could obtain for Barney and Betty in these addition cases.

Table 8. Unlimited Harm

	Barney	Betty
Option 1	0	Ω
Option 2	$x + 1$	$-x$
Option 3	0	y

For any $x, y > 0$, Option 2 is permissible according to Minimize Unanswered Complaints, regardless of whether its criteria of defeat imply that one option can defeat another when the latter causes non-identity shortfall relative to the former. In *Unlimited Harm*, none of the options causes non-identity shortfall. And since Barney exists given any option (i.e., he is not a contingent person), the harm to Betty done by Option 2 is wholly unanswered. Yet, no matter how awful Betty's life given Option 2 (i.e., no matter what negative value we assign to $-x$), and no matter how fabulous her life given Option 3 (i.e., no matter what positive value we assign to y), Option 2 remains uncovered, and therefore permissible. Neither Option 1 nor Option 3 can cover Option 2 on Minimize Unanswered Complaints. For any $x \geq 0$, Option 2 defeats Option 1, since the comparative harm that Option 1 does to Barney is greater than the existential harm that Option 2 does to Betty. For sufficiently large values of y in relation to x , Option 3 defeats Option 2. However, there are no values of x and $y \geq 0$, for which Option 3 defeats Option 1, since neither Option 1 nor Option 3 causes any harm or non-identity shortfall relative to the other.

Since x and y can take any values greater than 0, Option 2 can inflict any *greater* amount of permissible unanswered harm than either Option 1 or Option 3. Podgorski's Minimize Unanswered Complaints therefore violates the following principle, regardless of whether it is modified to accommodate NED:

Limit Permissible Harm (LPH): If option A does more unanswered harm than any alternative, and no alternative causes non-identity shortfall, then if A is permissible, the difference between the amount of unanswered harm done by A and that done by any alternative cannot be arbitrarily great.

In other words, there must be a limit to *how much more* unanswered harm a permissible option does relative to the alternatives.

Not only is LPH intuitively plausible, it is difficult to see how any harm-avoidance theory can reject it. According to the harm-avoidance account, the only possible wrong-makers for any of Options 1–3 in *Unlimited Harm* is the harm it does, since none of the options causes non-identity shortfall. How, then, can there be *no limit* to the amount of unanswered harm that is permitted? Podgorski's proposed criterion of permissibility, Uncovered, does not track what is morally relevant on the harm-avoidance account, namely *harm-avoidance*.

6. The End of the Road

So far, we have seen that three harm-avoidance theories violate Weaker DA, a fourth satisfies Weaker DA but violates WILA, while a fifth satisfies both Weaker DA and

WILA but violates LPH. This motivates the search for a harm-avoidance theory that accommodates all the aforementioned principles.

But we've come to the end of the road. No harm-avoidance theory can accommodate all three principles. Given two very weak assumptions, which we state below, the conjunction of these principles is incompatible with the defining feature of a harm-avoidance theory. Recall that according to Harmless Permission, an option that causes no existential harm, comparative harm, or non-identity shortfall is permissible. As we now demonstrate, Harmless Permission is incompatible with the conjunction of Weaker Dominance Addition, Weak Improvable Life Acceptance, and Limit Permissible Harm.

Consider Table 9.

Table 9. The End of the Road

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

The End of the Road is an abstract schema for a range of possible cases where x, y, z , and ε are well-being values for Persons 1 and 2, and these values can differ across different possible cases in the range.

The schema has five important features:

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

Feature 2: None of Options 1—3 causes non-identity shortfall, either in a binary choice or in a choice between all three Options.

Feature 3: For any $x, y, z, \varepsilon > 0$,

3a. In a choice between Option 1, Option 2, and Option 3, Person 1 is the only person harmed by Option 1.

3b. Option 3 harms Person 1 more than Option 1 does.

Feature 4: For any $x, y, z, \varepsilon > 0$, and $z - y > y - (x - \varepsilon)$, Option 3 does more unanswered harm than Option 1 or Option 2.

Feature 5: For any $x, y > 0$,

5a. Option 2 weakly addition-dominates Option 1.

5b. Everyone who exists given Option 1 has a bad life.

We adopt the following definitions:

Definition 1: Option A is morally required = *df.* A is permissible and any alternative to A is impermissible.

Definition 2: Option A is impermissible = *df.* A is not permissible.

Finally, we make the following two substantive but very weak assumptions:

Weak No Dilemma Assumption: For any $x, y, z, \varepsilon, > 0$, at least one of Options 1—3 is permissible.

Weak Completeness Assumption: For any $x, y, z, \varepsilon, > 0$, and for each one of Options 1—3, that option is either permissible or impermissible.

According to Weak No Dilemma Assumption, the cases that fit the schema *The End of the Road* where $x, y, z, \varepsilon > 0$ are not moral dilemmas. At least one of the options in those cases is permissible. This does not imply that *there are no moral dilemmas*. Hence, those who believe in the existence of moral dilemmas can accept Weak No Dilemma Assumption. But we think that *if* there are any moral dilemmas, there must be a special explanation as to why, in those choice contexts, every one of an agent's options is impermissible. We do not think that there is any such explanation to be given regarding the relevant cases that fit *The End of the Road*. The onus is on those who disagree to show why Weak No Dilemmas Assumption should be rejected.

According to Weak Completeness Assumption, in the cases that fit the schema *The End of the Road* where $x, y, z, \varepsilon > 0$, for each one of Options 1—3, its deontic status is either *permissible* or *impermissible*, not some third status, such as *indeterminate*. Again, we are not assuming that the deontic status of *any* option is either permissible or impermissible, only that this is true in the relevant range of cases. Like moral dilemmas, deontic indeterminacy is a phenomenon that requires special explanation, and we just don't see what the explanation could be in the cases that fit *The End of the Road*.

Given Weak No Dilemma Assumption, Weak Completeness Assumption, and Definitions 1 and 2, we can demonstrate that Harmless Permission, Weaker Domi-

nance Addition, Weak Improvable Life Acceptance, and Limit Permissible Harm are jointly incompatible.

Proof. Assume for reductio

P1. Harmless Permission: If option A does no comparative or existential harm, and does not create any non-identity shortfall, then A is permissible.

P2. Weak Improvable Life Acceptance: If (i) option A imposes greater harm on person S than option B, and (ii) the only moral consideration against B, in a choice from an option set \mathcal{O} that includes A and B, is that B harms S, then if A is not morally required in a binary choice between A and B, then A is not morally required in a choice from \mathcal{O} .

P3. Weaker Dominance Addition: If option A weakly addition-dominates option B, and everyone who exists given B has a bad life, then B is impermissible.

P4. Limit Permissible Harm: If option A does more unanswered harm than any alternative, and no alternative causes non-identity shortfall, then if A is permissible, the difference between the amount of unanswered harm done by A and that done by any alternative cannot be arbitrarily great.

From P1, Feature 1, and Definitions 1 and 2,

P5. For any $x, y, z, \varepsilon > 0$, in a binary choice between Option 1 and Option 3, Option 3 is not morally required.

From P1, P2, P5, and Features 2 and 3,

P6. For any $x, y, z, \varepsilon > 0$, in a choice between Option 1, Option 2, and Option 3, Option 3 is not morally required.⁴²

From P6, Weak No Dilemma Assumption, Weak Completeness Assumption, and Definitions 1 and 2,

⁴² Notice that P1 and Features 2 and 3 jointly imply that if Option 1 is impermissible in a choice between Options 1—3, then this can only be because Option 1 harms Person 1. In other words, the fact that Option 1 harms Person 1 is the only moral consideration against Option 1.

P7. For any $x, y, z, \varepsilon > 0$, and $z - y > y - (-x - \varepsilon)$, in a choice between Option 1, Option 2, and Option 3, either Option 1 is permissible or Option 2 is permissible.

From P4, and Features 2 and 4,

P8. For some $x, y, z, \varepsilon, > 0$, and $z - y > y - (-x - \varepsilon)$, in a choice between Option 1, Option 2, and Option 3, Option 2 is impermissible.

From P7, P8, and Definition 2,

P9. For some $x, y, z, \varepsilon, > 0$, and $z - y > y - (-x - \varepsilon)$, in a choice between Option 1, Option 2, and Option 3, Option 1 is permissible.

But from P3 and Feature 5,

P10. For any $x, y, z, \varepsilon, > 0$, in a choice between Option 1, Option 2, and Option 3, Option 1 is impermissible.

So, from P10, Definition 2, and existential instantiation,

C. It is not the case that for some $x, y, z, \varepsilon, > 0$, and $z - y > y - (-x - \varepsilon)$, in a choice between Option 1, Option 2, and Option 3, Option 1 is permissible.

Since C contradicts P9, we must reject either Harmless Permission, Weaker Dominance Addition, Weak Improvable Life Acceptance, or Limit Permissible Harm.

But Harmless Permission is part and parcel of any harm-avoidance theory. So it seems, proponents of a harm-avoidance theory must reject either Weaker Dominance Addition, Weak Improvable Life Acceptance, or Limit Permissible Harm. The worry is that each of these claims is extremely plausible, more plausible, we think, than Harmless Permission.

7. Conclusion

One of the central challenges facing any harm-avoidance theory, i.e., any theory committed to Harmless Permission, is offering an adequate response to *the Problem of Improvable Life Avoidance*. The problem is that the simplest harm-avoidance theory, Harm Minimization, leads to both improvable life avoidance and a requirement to choose weakly addition-dominated options. Most harm-avoidance theorists seem to

agree that improvable life avoidance and requiring weakly addition-dominated options are problematic. But in scrutinizing the existing harm-avoidance theories, and their responses to *the Problem of Improvable Life Avoidance*, we have argued that some of these theories permit weakly addition-dominated options in which everyone has a bad life, that some lead to strong improvable life avoidance, and that some permit any amount of harm in relation to less harmful alternatives. Moreover, we have argued that no harm-avoidance theory can avoid all three of these problems.

Our discussion bears on the prospects of finding a general theory that accommodates the Asymmetry. Harm-avoidance theories have seemed like the most promising candidates in this regard. In light of our discussion, one might be motivated to find an alternative theoretical framework in which to situate the Asymmetry. But the alternatives come with their own problems.

One possibility would be to defend the Asymmetry by appealing to a different type of harm-avoidance. For instance, one could claim that the only harm that we are morally required to avoid is non-comparative harm. We might be required to avoid causing people to be in an intrinsically bad state, but not to avoid giving people less of what is intrinsically good rather than more of what is intrinsically good.

However, this may seem quite extreme. It implies, for example, that we have no moral requirement to save people from death, insofar as death would not be intrinsically bad for those who die but would merely deprive them of further good.

An alternative response, which we find more plausible, is to reject the Asymmetry. We should accept that we *can* be morally required to create people with good lives rather than not create them at all, where the explanation for this is simply that these people would exist with good lives.

What about Climate Anti-Natalism, the claim that in many situations it is wrong to create a person because of the added CO₂ emissions? Insofar as Climate Anti-Natalism is motivated by the harm-avoidance account, our result undermines the case for Climate Anti-Natalism. In the domain of normative population ethics, there are several sources of an act being impermissible. It could be impermissible because it causes harm, because it causes non-identity shortfall, or because it fails to create people with good lives. We have argued against the claim that an act that causes no harm or non-identity shortfall in this domain is permissible. This makes it seem likely that there will be cases where an act is impermissible because it fails to create a person with a good life. Those who wish to defend Climate Anti-Natalism must therefore address the possible existence of such reasons, and show that they aren't strong enough to outweigh the expected climate-change-related harm of adding another person to the world.

References

- Arrhenius, G. (2003). The person-affecting restriction, comparativism, and the moral status of potential people. *Ethical Perspectives*, 10(3), 185–195.
- Arrhenius, G. (2022). Population paradoxes without transitivity. In G. Arrhenius, K. Bykvist, T. Campbell, & E. Finneron-Burns (Eds.), *The Oxford handbook of population ethics* (pp. 181–203). Oxford University Press.
- 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. (2022a). The asymmetry. In J. McMahan, T. Campbell, J. Gooddrich & K. Ramakrishnan (Eds.), *Ethics and existence: The legacy of Derek Parfit* (pp. 15– 37). Oxford University Press.
- Bader, R. (2022b). Person-affecting utilitarianism. In G. Arrhenius, K. Bykvist, T. Campbell & E. Finneron-Burns (Eds.), *The Oxford handbook of population ethics* (pp. 251–270). Oxford University Press.
- Benatar, D. (2006). *Better never to have been: The harm of coming into existence*. Oxford University Press.
- Boonin, D. (1996). Don't stop thinking about tomorrow: Two paradoxes about duties to future generations. *Philosophy & Public Affairs*, 25(4), 267–307.
- Boonin, D. (2014). *The non-identity problem & the ethics of future people*. Oxford University Press.
- Bradley, B. (2012). Doing away with harm. *Philosophy and Phenomenological Research*, 85(2), 390–412.
- Burkett, Daniel. (2021). A Legacy of Harm? Climate Change and the Carbon Cost of Procreation. *Journal of Applied Philosophy* 38: 790-808.
- 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.
- Conly, S. (2015). *One Child: Is there a right to more?* New York: Oxford University Press.
- Cusbert, J., & Kath, R. (2018). A consequentialist account of Narveson's dictum. *Philosophical Studies*, 176(7), 1693–1709.

- Frick, J. (2020). Conditional reasons and the procreation asymmetry. *Philosophical Perspectives*, 34(1), 53–87.
- Frick, J. (2022). Context-dependent betterness and the mere addition paradox. In J. McMahan, T. Campbell, J. Goodrich & K. Ramakrishnan (Eds.), *Ethics and existence: The legacy of Derek Parfit* (pp. 232–263). Oxford University Press.
- Gardner, M. (2015). A harm-based solution to the non-identity problem. *Ergo*, 2(17), 427–444.
- Gardner, M. (2019). David Boonin on the non-identity argument: Rejecting the second premise. *Law, Ethics and Philosophy*, 7, 29–47.
- 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.
- Hedberg, T. (2019). “The duty to reduce greenhouse gas emissions and limits of permissible procreation. *Essays in Philosophy* 20: 1–24.
- Herlitz, A. (2020). Non-transitive better than relations and rational choice. *Philosophia*, 48(1), 179–189.
- Holtug, N. (2010). *Persons, interests, and justice*. Oxford University Press.
- Horton, J. (2021). New and improvable lives. *Journal of Philosophy*, 118(9), 486–503.
- Kaczmarek, P. & Lloyd, H. R. (forthcoming). Moral uncertainty, pure justifiers, and agent-centred options. *Australasian Journal of Philosophy*, 00, 1–29. [https : //philpapers.org/rec/KACMUP-2](https://philpapers.org/rec/KACMUP-2).
- Kagan, S. (1991). Replies to my critics. *Philosophy and Phenomenological Research*, 51(4), 919–928.
- Kamm, F. M. (1985). Supererogation and obligation. *Journal of Philosophy*, 82(3), 118–138.
- 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.

- MacIver, C. (2015). Procreation or appropriation? In Sarah Hannan, Samantha Brennan & Richard Vernon, eds. *Permissible Progeny?: The Morality of Procreating and Parenting*, 107-128. Oxford University Press.
- McDermott, M. (1982). Utility and population. *Philosophical Studies*, 42(2), 163–177.
- 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-futuredeontic-long-termism-and-the-non-identity-problem/>.
- Mogensen, A. (2021). Moral demands and the far future. *Philosophy and Phenomenological Research*, 103(3), 567–585.
- Muñoz, D. (2021). From rights to prerogatives. *Philosophy and Phenomenological Research*, 102(3), 608–623.
- Otsuka, M. (1997). Kamm on the morality of killing. *Ethics*, 108(1), 197–207.
- 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: Oxford University Press.
- Parfit, D. (2017). Future people, the non-identity problem, and person-affecting principles. *Philosophy & Public Affairs*, 45(2), 118–157.
- Persson, I. (2017). *Inclusive ethics: Extending beneficence and egalitarian justice*. Oxford University Press.
- 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.
- 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.
- Schwartz, T. (1972). Rationality and the myth of the maximum. *Noûs*, 6(2), 97–117.
- Shiffrin, S. (1999). Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory*, 5, 117–148.
- 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.
- Wynes, S. and Nicholas, K. (2017). The climate mitigation gap: Education and government recommendations miss the most effective individual actions. *Environmental Research Letters* 12: 3.
- Young, T. (2001). Overconsumption and procreation: Are they morally equivalent? *Journal of Applied Philosophy* 18 (2001): 183–192.

Appendix. Minimize Unanswered Complaints Satisfies Weak Dominance Addition

We will show that on Podgorski's Minimize Unanswered Complaints:

Claim: For any choice context in which A and B are both options, if A weakly addition dominates B, then A defeats B and for any C, if B defeats C, then A defeats C.

It follows from Claim that if A weakly addition dominates B, then A covers B, and thus B is impermissible on the Uncovered criterion. So, the conditional 'if B is permissible, then A is permissible' is vacuously true.

Proof. Suppose that in some choice context,

P1. A weakly addition dominates B.

First we will prove that given P1, A defeats B. From P1 and the definition of 'weak addition dominance',

P2. Every person who exists in B exists in A, and every person who exists in B has well-being at least 0 and at most x , and every person who exists in A has positive well-being $y > x$.

Let $\text{Harm}(A_B)$ represent the total harm in A relative to B. Because total harm is the sum of existential harm and comparative harm, from P2 and the definitions of 'existential harm', 'comparative harm', and 'existential benefit answers' we derive P3—P5:

P3. $\text{Harm}(A_B) = 0$.

P4. $\text{Harm}(B_A) > 0$.

P5. $\text{ExBen}(B_A) = 0$.

From P4 and P5, we derive

P6. $\text{Harm}(B_A) - \text{ExBen}(B_A) > 0$

From P3 and P6, we derive

P7. $\text{Harm}(B_A) - \text{ExBen}(B_A) > \text{Harm}(A_B) - \text{ExBen}(A_B)$.

From P6, P7, and Minimize Aggregate Unanswered Harm's definition of 'defeat', i.e., MAUH, we derive

P8. A defeats B.

Next, we will prove that for any alternative C in the choice context, if B defeats C, then A defeats C. Suppose:

P9. There is some C in the choice context, such that B defeats C.

Since B defeats C, by the definition of 'defeat',

P10. $\text{Harm}(C_B) - \text{ExBen}(C_B) > 0$

and

P11. $\text{Harm}(C_B) - \text{ExBen}(C_B) > \text{Harm}(B_C) - \text{ExBen}(B_C)$.

From P10, we derive

P12. $\text{Harm}(C_B) > 0$.

We will show that for any individual harm in C relative to B, there is at least that much individual harm in C relative to A.

First, any harm in C relative to B is either existential harm in C relative to B or comparative harm in C relative to B.

Any existential harm in C relative to B is suffered either by someone who exists in C but neither A nor B, or by someone who exists in C and A but not in B. Any person who suffers existential harm in C relative to B, has negative well-being $-z$ in C. If the person exists in C but neither A nor B, then the magnitude of her existential harm in C relative to B and that of her existential harm in C relative to A is $|-z|$. If she exists in C and A but not B, then the harm she suffers in C relative to A is comparative, and the magnitude of this harm in C relative to A is $|-z|+y$, i.e. the difference between her positive welfare in A and her welfare in C, where $(|-z| + y) > |-z|$. Hence, for any existential harm in C relative to B, there is either that much existential harm in C relative to A, or even greater comparative harm in C relative to A.

Next, any comparative harm in C relative to B is suffered by someone who exists in both B and C. Suppose she has lifetime well-being w_C in C, where w_C could be any positive or negative number, $w_C < x$. Since everyone who exists in B exists in A with welfare y , and $y > x$, it straightforwardly follows that $w_C < x < y$. Hence, any

comparative harm in C relative to B is an even greater comparative harm in C relative to A.

It follows that for any individual harm in C relative to B, there is at least that much individual harm in C relative to A. Since the total harm in one option relative to another is just the sum of the individual harms in the former relative to the latter, we derive

$$\mathbf{P13.} \text{Harm}(C_A) \geq \text{Harm}(C_B).$$

From P10 and P13, we derive

$$\mathbf{P14.} \text{ExBen}(C_B) < \text{Harm}(C_A).$$

By the definition of ‘existential benefits’, any existential benefit in C relative to B is had by a person who exists in C but not B. Hence, any existential benefit in C relative to B is had either by a person who exists in C but not in A or B, or by a person who exists in C and A, but not B. Hence,

$$\mathbf{P15.} \text{ExBen}(C_B) = \text{the sum of positive well-being of (i) all people who exist in C but not in A or B and (ii) all people who exist in C and A but not B.}$$

Let W_C = the sum of positive well-being of all people who exist in C but not A or B. And let W_{CA} = the sum of positive well-being of all people who exist in C and A but not B.

Then, from P14 and P15, we derive

$$\mathbf{P16.} W_C + W_{CA} < \text{Harm}(C_A).$$

From P16, we derive

$$\mathbf{P17.} W_{CA} < \text{Harm}(C_A) \text{ and}$$

$$\mathbf{P18.} W_C < \text{Harm}(C_A).$$

The existential benefits of C relative to A consists of the positive well-being of those who exist in C but not A. Since everyone who exists in B exists in A, this means that the existential benefits of C relative to A consists of the positive well-being of those who exist in C but neither A nor B, i.e.,

$$\mathbf{P19.} \text{ExBen}(C_A) = W_C.$$

From P18 and P19, we derive

$$\mathbf{P20.} \text{ExBen}(C_A) < \text{Harm}(C_A).$$

Subtracting $\text{ExBen}(C_A)$ from both sides of the inequality in P20, we derive

$$\mathbf{P21.} \text{Harm}(C_A) - \text{ExBen}(C_A) > 0$$

Next, we need to see how $\text{Harm}(A_C)$ compares to $\text{Harm}(C_A)$.

There are no existential harms in A relative to any outcome. Hence, if there is any individual harm in A relative to C, it is a comparative harm, i.e., a harm to someone who exists in A and C. Everyone who exists in A has positive well-being y . By the definition of ‘comparative harm’, for anyone harmed in A relative to C, her well-being in C is greater than her well-being in A, (i.e., $> y$). Call those harmed in A relative to C ‘the A-harmed people’. Let $W_{A\text{-harmed}}^A$ be the sum total of positive well-being of the A-harmed people in A. Let $W_{A\text{-harmed}}^C$ be the sum total of positive well-being of the A-harmed people in C. The total harm of A relative to C is therefore equal to $W_{A\text{-harmed}}^C - W_{A\text{-harmed}}^A$. In other words,

$$\mathbf{P22.} \text{Harm}(A_C) = W_{A\text{-harmed}}^C - W_{A\text{-harmed}}^A.$$

Since W_{CA} is the total positive well-being of everyone who exists in both A and C,

$$\mathbf{P23.} W_{A\text{-harmed}}^C \leq W_{CA}.$$

Since $W_{A\text{-harmed}}^A$ is a positive number, from P23, we derive

$$\mathbf{P24.} W_{A\text{-harmed}}^C - W_{A\text{-harmed}}^A < W_{CA}.$$

From P22 and P24, we derive

$$\mathbf{P25.} \text{Harm}(A_C) < W_{CA}.$$

From P16, P19, and P25, we derive

$$\mathbf{P26.} \text{ExBen}(C_A) + \text{Harm}(A_C) < \text{Harm}(C_A).$$

Subtracting $\text{ExBen}(C_A)$ from both sides of the inequality in P26, we get:

$$\mathbf{P27.} \text{Harm}(A_C) < \text{Harm}(C_A) - \text{ExBen}(C_A).$$

The unanswered harm of C relative to A is greater than the harm of A relative to C, and hence, greater than the unanswered harm of A relative to C. Finally, from P21, P27, and the definition of 'defeat', C. A defeats C

We have proven that if A weakly addition dominates B, then A defeats B and for any C, if B defeats C, A defeats C. From which it follows that if A weakly addition dominates B, A covers B. Therefore, Podgorski's Minimize Unanswered Complaints satisfies Weak Dominance Addition.