

A FISSION PROBLEM FOR PERSON-AFFECTING VIEWS

ELLIOTT THORNLEY

Massachusetts Institute of Technology

On person-affecting views in population ethics, the moral import of a person's welfare depends on that person's temporal or modal status. These views typically imply that—all else equal—we're never required to create extra people, or to act in ways that increase the probability of extra people coming into existence.

In this paper, I use Parfit-style fission cases to construct a dilemma for person-affecting views: either they forfeit their seeming-advantages and face fission analogues of the problems faced by their rival impersonal views, or else they turn out to be not so person-affecting after all. In light of this dilemma, the attractions of person-affecting views largely evaporate. What remains are the problems unique to them.

1. Introduction

Suppose that you find yourself with a choice. You can either:

(a) Donate \$4500 to the Against Malaria Foundation (AMF).

Or:

(b) Donate \$4500 to the Nuclear Threat Initiative (NTI).

You're confident that donating to AMF would save a child from dying of malaria. You're also reasonably sure that this child would go on to live an additional 70 years of good life. On the other hand, you estimate that donating to NTI would increase the probability that humanity survives the coming century by about one-in-ten-quadrillion (10^{-16}). And you expect that if humanity survives the coming century, the future will contain one-hundred-quadrillion (10^{17}) good lives, each lasting around 70 years. Where should you send your money?

Contact: Elliott Thornley <elliott.thornley@gmail.com>

Here's a quick argument for NTI. By donating to AMF, you'd cause about 70 additional years of good life to be lived, in expectation. By donating to NTI, you'd cause about 700 additional years of good life to be lived, in expectation. It's better to add 700 years of good life than it is to add 70 years of good life. Therefore, you should send your money to NTI.

There are many ways to resist this quick argument, but perhaps the most natural way is to claim that the years of good life that might result from your NTI donation *just don't matter* in the same way as the years of good life that would result from your AMF donation. By donating to AMF, you gift 70 more years to a person who *actually* exists, who *will* exist regardless of your decision, and who exists *right now*. The same can't be said of your donation to NTI. The vast majority of those additional years would accrue far in the future: to people who do not and need never exist.

This is a *person-affecting* response to the quick argument. On person-affecting views in population ethics, the moral import of a person's welfare depends on that person's temporal or modal status. These views typically imply that—all else equal—we're never required to create extra people, or to act in ways that increase the probability of extra people coming into existence.

The allure of person-affecting views is partly in their foundations. These views often have their start in two claims that many find intuitive: (1) the *Person-Affecting Restriction*: an outcome can't be better than another unless it's better for some person, and (2) *Existence Anticomparativism*: existing can't be better for a person than not existing.

However, another big draw of person-affecting views is their upshots. These views avoid some well-known problems faced by their rival impersonal views. Consider expected total utilitarianism: one prominent impersonal view. It implies that there are cases in which we're required to create new happy people rather than help existing people, cases in which we're required to make great sacrifices to create new people with lives barely worth living, and cases in which we're required to make great sacrifices to slightly reduce the chance of near-term human extinction. Person-affecting views mostly avoid these problems, and that might seem like a significant point in their favour.

In this paper, I argue that these advantages are largely illusory. Using Parfit-style fission cases (Parfit 1984: §89), I construct a dilemma for person-affecting views: either these views violate the spirit of the Person-Affecting Restriction, or else they imply fission analogues of the problems that blight impersonal views. These fission analogues are about as troubling as the original problems, and so they undermine much of the motivation for preferring person-affecting views to impersonal views. Considering the objections unique to person-affecting views, we should prefer impersonal views on balance.

Rejecting person-affecting views doesn't immediately commit us to NTI over AMF. There are many ways to resist the quick argument. But—as I hope to show in this paper—the most natural line of resistance isn't as attractive as it might first seem.¹

2. Person-Affecting Views

I've defined 'person-affecting views' as follows: those views on which the moral import of a person's welfare depends on that person's temporal or modal status.² Such views typically designate some people as *extra*, and then claim that the welfare of these extra people doesn't matter in the same way as the welfare of non-extra people. On *presentism*, it's future people that are extra. On *actualism*, it's non-actual people: those who don't and won't exist in the actual world. On *necessitarianism*, it's non-necessary people: those whose

1. In a companion paper (Thornley 2025), I argue that fission also presents a challenge to critical-level and critical-range views in population ethics. In that paper's introduction, I give a brief argument against such views, intended to save the time of readers of a certain metaphysical bent. Here's the analogous argument against person-affecting views:

1. On person-affecting views, our moral obligations can depend on the affected persons' temporal or modal status.
2. A person's temporal or modal status can depend on our answers to questions of personal identity. (Whether a person *presently*, *actually*, or *necessarily* exists in some scenario—or whether they're *harmed* by some action—can depend on whether that person is identical to some person existing at other times or in other possible worlds.)
 - C1. So, on person-affecting views, our moral obligations can depend on our answers to questions of personal identity.
3. Questions of personal identity are *empty*: their answers can't be discovered but at most stipulated.
4. Our moral obligations can't depend on an answer to an empty question.
 - C2. Therefore, person-affecting views are false.

I have some sympathy for this argument, but my case against person-affecting views doesn't depend on it.

2. There are a couple of complexities to note here. First, Bader's (2022) *same-number utilitarianism* doesn't discriminate on temporal or modal grounds, but counts as person-affecting on another natural definition of the term: The view implies that it's never better to create new people, all else equal. On Bader's view, populations of the same size are ordered by sum-totals of welfare, while populations of different sizes are incomparable. I explain how my arguments apply to Bader's view in fn 22.

Second, some define 'person-affecting views' as all and only those views that satisfy the Person-Affecting Restriction. That would make total utilitarianism paired with the negation of Existence Anticomparativism a person-affecting view. It would also imply that *wide views* (explained below) paired with Existence Anticomparativism are not person-affecting. Since my arguments tell against wide views but not total utilitarianism, I use the definition of 'person-affecting views' to which this footnote is appended.

existence depends on our choice. On *comparativism*, it's people who exist in just one of two compared outcomes.³

Harm-minimisation views (HMs) are a slightly different matter. As the name indicates, they ask us to minimise *harm*, understood as the amount by which a person's welfare falls short of what it could have been. What makes these views paradigmatically person-affecting is their claim that a person can't be harmed in an outcome in which they don't exist.⁴ HMs don't categorise people as extra and non-extra *simpliciter*, but we can understand them to designate people as extra *in an outcome A relative to an outcome B*. A person is extra in this way iff that person exists in *A* but not in *B*. In the two-option cases I discuss below, I often write that people are extra *simpliciter*. Applied to HMs, I mean that they are extra *in the outcome in which they exist, relative to the other available outcome*.

Each of the above five classes of person-affecting view is broad. As stated, they leave many issues unsettled. One issue is how to treat the welfare of extra people living bad lives. On *symmetric* views, the welfare of extra people doesn't matter at all, whether their lives are good or bad. Many find symmetric views implausible, due in part to cases like the following. Suppose that Nikita imposes some small cost on non-extra people to prevent the creation of a huge number of extra people living awful lives. If extra people's welfare doesn't matter at all (and there are no relevant non-welfarist considerations in play), Nikita's act is wrong. But her act seems right.⁵ That intuition might lead us to prefer an *asymmetric* view, on which the welfare of extra people living bad lives matters in the same way as the welfare of non-extra people, while the welfare of extra people living good lives does not.

Here's a second dimension along which person-affecting views can vary. They can be *soft*, *hard*, or *very hard*, depending on the way in which they take extra good lives to matter.⁶ To see the difference, consider the following populations:

Soft, Hard, or Very Hard

Population A		Population B		Population C		Population D	
Nicholas	100	Nicholas	100	Nicholas	100	Nicholas	99
Vivianne	Ω	Vivianne	g	Vivianne	g	Vivianne	g
Mana	Ω	Mana	Ω	Mana	-1	Mana	Ω

3. As stated, comparativism applies only in two-option cases. The view is usually supplemented with a rule that determines what we're permitted to do in cases with three-or-more options (Ross 2015: §5; Thomas 2023: §4.1).

4. Or, on Roberts' (2011b: 356) view: any harms to a person are morally insignificant in outcomes in which they don't exist.

5. This case is a generalisation of Hare's (2007: 499) "Childless George" case.

6. Here and below, I use 'extra good lives' as shorthand for 'the welfare of extra people living good lives.' The same goes for my use of 'extra bad lives.' The 'soft' and 'hard' labels come from Thomas (2023: 487).

The numbers in this table represent people’s welfare. Positive numbers indicate good lives and negative numbers indicate bad lives. Ω indicates that a person doesn’t exist in a population.

Population B is identical to population A but for the addition of Vivianne, living a good life with welfare g . Population C adds Vivianne too, but this time at some cost: Mana lives a bad life in C, while in A she lives no life at all. Population D also adds Vivianne at some cost: Nicholas is worse off in D than he is in A.

As noted above, person-affecting views typically imply that—all else equal—we’re never *required* to create extra good lives. That means that no matter how good Vivianne’s life is—no matter how large g is—we’re never required to choose B over A. Either is permissible in a choice between the two.

On very hard views, we’re also never *permitted* to create extra good lives if doing so involves the creation of extra bad lives or any harm to non-extra people. That means that, no matter how good Vivianne’s life is, we’re never permitted to choose C over A, or D over A.

Hard views also forbid creating extra good lives if doing so harms non-extra people, but they permit creating combinations of extra good and bad lives, so long as the good lives are good enough. That means that we’re required to choose A over D, but permitted to choose C over A for large enough values of g .

On soft views, by contrast, creating extra good lives can be permissible both when doing so involves creating extra bad lives and when doing so harms non-extra people. So long as Vivianne’s life is good enough, we’re permitted to choose C over A, and D over A.

Here’s a third dimension along which species of actualism, necessitarianism, comparativism, and HMs can differ.⁷ Such views can be *narrow* or *wide*.⁸ To see the difference, consider the following *Non-Identity Case* (Parfit, 1984: ch. 16):

Non-Identity Case

Population D		Population E	
Healthy	100	Healthy	Ω
Unhealthy	Ω	Unhealthy	1

In this case, narrow views permit us to create either Healthy or Unhealthy. That’s because narrow views are defined as those views that use transworld identity as their counterpart relation for the purposes of determining which persons are extra. Healthy and Unhealthy are not identical, so both count as extra on narrow necessitarianism, comparativism, and HMs, and Healthy counts as extra

7. The distinction concerns how a person’s modal status is determined, and so doesn’t apply to presentism.

8. These labels are also borrowed from Thomas (2023: 490). I note, as he does, that they’re a close but imperfect match for traditional terminology.

if we create Unhealthy (and vice versa) on narrow actualism. Since Healthy and Unhealthy are both extra, we're granted broad latitude in choosing who to create. Wide views, on the other hand, require us to create Healthy. That's because wide views are defined as those views that employ counterpart relations that extend transworld identity. These extended counterpart relations first pair people up by identity, and then go on to pair up some non-identical people. The relations offered in the literature are typically *saturating*—they pair up as many people as possible—and so imply that Healthy and Unhealthy are counterparts (Meacham 2012: 266–267; Thomas 2016: 211; 2023: 30–31). On wide views, therefore, both Healthy and Unhealthy count as non-extra, and their welfare matters accordingly. Plausible views will require that we bestow larger rather than smaller benefits on non-extra people, and so require that we create Healthy.

The above dimensions give some sense of the variety of possible person-affecting views. Even so, many views in the literature don't slot neatly into the resulting taxonomy. That's partly because the taxonomy doesn't map the entirety of logical space and partly because many person-affecting views are sketched out in strokes too broad to determine where they fall along certain axes.⁹

3. Seeming-Advantages of Person-Affecting Views

While person-affecting views vary widely in their details, they're largely united in their attractions. As noted above, many person-affecting views are founded on

9. These include Kamm's (2005: 304–305) view, which seems largely presentist, Narveson's (1973: 65) and Warren's (1977: 285) views, which seem largely actualist, Heyd's (1988) view, which seems presentist in some places and necessitarian at others, Heyd's (1992: 97) view, which seems necessitarian in some places and actualist at others, Setiya's (2014) view, which seems actualist in some places and presentist at others, and Bigelow's and Pargetter's view (1988), which seems presentist, necessitarian, and actualist at different points. Ross (2015) offers a comparativist view, but suggests that our obligations also depend on non-person-affecting considerations. Thomas (2023) constructs four views—each asymmetric and comparativist—filling a 2×2 grid of soft/hard and narrow/wide. Singer (2011: 88–90) and Bradley (2013) both discuss—but do not endorse—an asymmetric, necessitarian view. Parsons (2002) suggests a symmetric, actualist view that seems more asymmetric in its deontic upshots. Cohen's (2020) "Subjective Actualism" is an asymmetric, very hard actualist view that's narrow in canonical non-identity cases but wide in more realistic cases. Spencer's (2021) "Stable Actualism" is an asymmetric, narrow form of actualism. Hare (2007) offers a wide form of actualism. McDermott (1982) constructs an asymmetric, narrow, very hard HMV. Roberts' (2011b) "Variabilism" is also an asymmetric, narrow HMV, albeit with the caveat that her view states only which harms are morally significant. It doesn't state how these harms bear on our moral obligations. Temkin (2012: ch. 12) seems to lean towards a narrow HMV, though like Ross (2015) he suggests that our obligations also depend on non-person-affecting considerations. Meacham's (2012) "Saturating Harm Minimizing View" is an asymmetric, wide HMV. Mogensen's (2019) "Non-Requiring View+" is asymmetric, narrow, and soft, as is Horton's (2021) "Avoid Reasonable Objections" view and Podgorski's (2023) view. Pummer's (2024) view is asymmetric and soft. McDermott's (2019) "Objection Minimization" view is asymmetric, narrow, and very hard.

two *prima facie* appealing claims: (1) the Person-Affecting Restriction: an outcome can't be better than another unless it's better for some person, and (2) Existence Anticomparativism: Existing can't be better for a person than not existing.¹⁰ Theories that violate the Person-Affecting Restriction can seem objectionably impersonal: treating people as mere containers of value (Bader 2022: 2–4; Frick 2017: 351; Holtug 2004: 131–132; Nebel 2021: 9, 12–13; Parfit 1984: 393–394; though see Chappell 2015: §3.1 for a response). Denying Existence Anticomparativism, meanwhile, seems to land us in a metaphysical tangle: if existing is better for a person than not existing, then it seemingly must be that not existing would be worse for that person than existing. But how can anything be better or worse for a person that doesn't exist?¹¹

Person-affecting views also have some attractive upshots. Extant views imply something in the vicinity of Narveson's Slogan: "We are in favor of making people happy, but neutral about making happy people" (1973: 80). Many views also imply what Roberts (2011a: 772) and Chappell (2017: 170) call the *Deeper Intuition*: we ought to benefit an existing person by some amount *g* rather than create a new person with welfare *g*. The exceptions are soft views, which may permit us to create the new person in this case. However, even soft views never *require* that we create the new person, even if that person's welfare would be much greater than *g*.¹² That's another implication which many find appealing. There might seem to be something perverse about theories that could require us to create new lives rather than help those suffering today.

Person-affecting views also avoid an especially pernicious version of Parfit's Repugnant Conclusion (1984: 388), which we can call *Repugnant Transition*. Suppose that everyone on earth is set to live a wonderful life. Suppose also that we

10. I write 'many' because not all person-affecting views uphold these claims. Roberts (2011b: 338) denies Existence Anticomparativism, and wide views paired with Existence Anticomparativism are tough to square with the Person-Affecting Restriction: In our Non-Identity Case, creating Healthy is required even though it's not better for anyone than creating Unhealthy.

11. Broome (1999: 168) gives an argument along these lines. Greaves and Cusbert (2022) argue that it fails.

Although the Person-Affecting Restriction and Existence Anticomparativism each have their charms, there are some difficulties associated with their conjunction. Together they imply that creating Unhealthy is no worse than creating Healthy in our Non-Identity Case above. Some find that verdict counterintuitive (Parfit 1984: §123). And an analogue of Broome's argument for Existence Anticomparativism implies that existing cannot be *worse* for a person than not existing. Coupled with the Person-Affecting Restriction, that claim entails that creating a person with an awful life is no worse than creating no one at all. Many find that implication troubling.

12. One caveat: depending on the new person's welfare and on how non-extra welfare is aggregated, *strong* actualism might imply that *if we create the new person* we're required to create the new person. However, this requirement won't have the usual force from the *ex ante* perspective, when we're deciding what to do. That's because strong actualism also implies that *if we don't create the new person* we're required not to create them. For more details on the distinction between strong and weak actualism, see Hare (2007).

could burden ourselves to such an extent that our lives would only be barely worth living, while creating many extra lives that are also barely worth living. On *total utilitarianism* and some other impersonal views, we're required to do so if the number of extra lives is large enough.¹³ On person-affecting views, we face no such requirement. On soft views, we're at most permitted to make the transition, while hard and very hard views forbid it.¹⁴

Another attraction of person-affecting views is their implications in more realistic cases. It's increasingly recognised that humanity's future hangs in the balance (Greaves & MacAskill 2025; Ord 2020). Here's one way it could play out. Earth supports a population of ten billion people per century until it becomes uninhabitable: one billion years from now. Future people do away with the sources of present-day suffering and cultivate much more of all that makes life good. As a consequence, Earth plays host to one-hundred-quadrillion (10^{17}) wonderful lives. Call this the *Good Future*. Here's another possible story. Runaway climate change, nuclear war, the release of an engineered pathogen, or some other disaster causes humanity to go extinct a century from now, soon after the lives of the present generation have run their course. Call this the *Short Future*.

There currently exist around eight billion people on Earth. Suppose for the sake of argument that we're all on course to live wonderful lives. Suppose also that we—the present generation—can shift the probabilities with which the Good and Short Futures come about. By all worsening our lives so that they're just barely worth living, we can decrease the Short Future's probability by p and increase the Good Future's probability by p . The other option is business-as-usual. For what values of p must we worsen our lives? On *expected total utilitarianism*, the answer is roughly "Any value greater than or equal to 0.0000008."¹⁵ We're required to make enormous sacrifices for the sake of people that may never exist, even if those sacrifices have just an eight-in-ten-million probability of paying off. Call this implication *Our Sacrifice*. Person-affecting views avoid this implication. It remains an open question how person-affecting views should be extended to cover risky cases (see Thomas 2023). But even in the case where $p = 1$, where business-as-usual would guarantee the Short Future and our generation's sacrifice would guarantee the Good Future, hard

13. Total utilitarianism states that a population's value is the sum-total of welfare in that population, and that bringing about a population is permissible iff no other available population has greater value.

14. Strong actualism might require the transition *if we make the transition*. See fn 12.

15. Expected total utilitarianism is the conjunction of total utilitarianism and expected value theory. On this view, downgrading eight billion lives from wonderful to barely-worth-living is *almost* as bad as *removing* eight billion wonderful lives, but increasing the chance of the Good Future by 0.0000008 is *as good* as *creating* eight billion wonderful lives. So, at $p = 0.0000008$, the benefits of present-day-sacrifice outweigh the costs. This figure is only rough, in part because my calculation ignores the welfare of the small number of future people in the Short Future.

and very hard views forbid the sacrifice since it harms us non-extra people. Soft views at best permit it.

So person-affecting views mostly avoid the problems above, and that might seem like a significant advantage of these views. In the next section, I'll argue that this advantage is largely illusory. Any person-affecting view worthy of the name will imply fission analogues of the problems for its rival impersonal view. These fission analogues are about as troubling as the original problems, and so the original problems give us little reason to prefer person-affecting views.¹⁶

What's more, we have more than little reason for the opposite preference. Person-affecting views face objections that impersonal views do not. Here's a brief survey. Symmetric person-affecting views are implausibly permissive about creating people with bad lives, while asymmetric person-affecting views must find a way to justify their asymmetry (Heyd 1988: 157, 161).¹⁷ Narrow person-affecting views imply — counterintuitively to many — that we're permitted to create the worse-off person in non-identity cases (Parfit 1984: ch. 16), and variations of Parfit's non-identity case pose problems for both narrow and wide views (Thornley 2024). Francis has argued that creating happy people is good (ms) and that creating happier people is better (2021); these arguments can't be defused by appeal to the Person-Affecting Restriction or Existence Anticomparativism, so person-affecting views need some other response. And a referee for this journal points out that person-affecting views can deliver implausible verdicts in a variation of Our Sacrifice. Suppose that the present generation will live wonderful lives no matter what, but that there will be no future generations unless I endure the small pain of scratching my finger. Bracketing off any non-welfarist considerations (though see Frick 2017), hard views require (and soft views permit) that my finger remains unscathed, no matter how numerous and wonderful the future lives at stake.¹⁸ That seems counterintuitive. Hume may be right that it's "not contrary to reason to prefer the destruction of the whole world to the scratching of my finger" (1738: 2.3.3.6), but it does seem contrary to morality.¹⁹

All this to say, we have good reason to favour impersonal views over person-affecting views. I now argue that we have little reason for the opposite preference.

16. It's also worth noting that we can avoid counterintuitive implications like Repugnant Transition and Our Sacrifice without moving to a person-affecting view. For example, we can reject *maximising consequentialism*: the claim that we're always required to do what's best.

17. For some proposed justifications, see Roberts (2011a; 2011b), Nebel (2019b), Frick (2020), Podgorski (2023), and Pummer (2024).

18. As above, strong actualism is an exception: I'm required to scratch my finger *if I scratch my finger*. See fn 12.

19. For other objections to person-affecting views, see Arrhenius (forthcoming: ch. 10), Beckstead (2013: ch. 4), Greaves (2017), Horton (2021), Ross (2015), Thomas (2023), and Thornley (2023).

4. Fission

Let a *life-episode* be a period within a person's life, and assume that each life-episode's welfare can be represented by a real-valued function w , such that life-episode x has at least as much welfare as life-episode y iff $w(x) \geq w(y)$. Assume also that welfare is interpersonally comparable (so that we can say whether x has at least as much welfare as y even if x and y are lived by different people) and measurable on a ratio-scale (so that we can talk meaningfully about the ratios of welfare between life-episodes). Assign positive welfare to life-episodes that are good for a person to live, negative welfare to life-episodes that are bad for a person to live, and zero welfare to life-episodes that are neither good nor bad for a person to live.

Suppose that we have the chance to split Anna's brain in two, and implant each half into an exact replica of her body. Each of the resulting people (call them Lefty and Righty) would share all of Anna's psychological features. Each of Lefty and Righty would also be phenomenally, physically, and functionally continuous with pre-fission Anna. That is to say, Anna's stream of (and capacity for) consciousness would divide and flow uninterrupted into the streams of (and capacities for) consciousness of Lefty and Righty.²⁰

If we choose No Split, Anna will live a life of welfare 80 and then die. If we choose Split A, Anna will live a life-episode of welfare 70 before the split. Both Lefty and Righty will then live *life-branches* of welfare 100. By a 'life-branch,' I mean a life-episode that begins immediately post-fission and ends with either fission or death.

Fission

No Split		Split A	
Anna	80	Anna	70
Lefty	Ω	Lefty	100
Righty	Ω	Righty	100

Suppose that we opt for No Split. In that case, which if any of Lefty and Righty should person-affecting views designate as extra? Here are six possible answers:

- (1) Each of Lefty and Righty is extra.
- (2) Lefty is extra.
- (3) Righty is extra.
- (4) Each of Lefty and Righty is "half-extra."
- (5) One of Lefty and Righty is extra, but it is indeterminate which.
- (6) Neither Lefty nor Righty is extra.

20. This case is a cosmetic variation on Parfit's *My Division* (1984: 254–255).

Take (1) first. If each of Lefty and Righty is extra, hard and very hard views imply that we were required to choose No Split, while soft views imply that choosing No Split was at least permitted. The nominal justification is that the only non-extra person—Anna—fares better in No Split. But these verdicts seem implausible, and are in fact hard to square with the Person-Affecting Restriction. That’s because—contrary to the above—Anna seems to fare better in Split A.

At least two lines of argument support this claim. The first is that Anna’s relation to Lefty and Righty seems to contain everything that could possibly matter in survival: she’s physically, psychologically, phenomenally, and functionally connected to both. The second is a two-step argument from Parfit (1984: 261–262). Start by imagining an outcome like Split A but with the right half of Anna’s brain destroyed, so that only Lefty exists. That seems better for Anna than No Split, since Lefty’s life-branch is wonderful and Anna is continuous with Lefty in all of the ways that might matter. Then reintroduce Righty, and note that it’s hard to see how this could make Anna worse off. She’s now continuous-in-all-the-ways-that-might-matter with two humans living wonderful life-branches rather than one, and “[h]ow could a double success be a failure?” (256).

No Split isn’t better than Split A for Lefty or Righty: they live wonderful life-branches in Split A and no life at all in No Split. If (as the above arguments suggest) No Split isn’t better for Anna either, then the Person-Affecting Restriction implies that No Split isn’t better overall. Hard and very hard views paired with (1) then seem objectionably impersonal, since they imply that we were required to choose No Split over Split A. And although soft views paired with (1) don’t violate the letter of the Person-Affecting Restriction in this case, they do seem to violate its spirit. The last paragraph’s arguments suggest that Split A is better than No Split for the only non-extra person: Anna.²¹ And the extra people in Split A—Lefty and Righty—both live wonderful life-branches. Given these facts, it seems that any person-affecting view worth the name would require you to choose Split A.²²

21. What Holtug (2010) calls *the Additive Prudential View* has a similar implication. On this view, “A person’s self-interest is an additive function of the (appropriately discounted) benefits that befall persons to whom she is M-related” (2010: 118), where the M-relation is “The continuous physical realization of a core psychology, and/or a distinctive psychology, and/or a chain of distinctive psychology” (2010: 99). This view implies that splitting is in Anna’s self-interest.

22. In fn 2, I promised to explain how my argument applies to Bader’s (2022) same-number utilitarianism. I now make good on that promise. Since (pace Dainton 1992) Anna, Lefty, and Righty are not all identical to each other (see fn 29), the population in Split A is larger than the population in No Split. Same-number utilitarianism thus implies that the two populations are incomparable. But—I’ve argued—Split A is wonderful for Lefty and Righty, and better for Anna than No Split. The spirit of the Person-Affecting Restriction thus requires that Split A is better than No Split, contrary to the verdict of same-number utilitarianism.

Answer (1), then, seems untenable. How about answers (2) and (3)? Perhaps person-affecting views should designate just Lefty as extra or just Righty as extra. But on reflection these answers also seem untenable. The left half of Anna’s brain could be identical to the right half in all relevant respects, and Lefty and Righty could start their life-branches sharing all relevant features. In that case, there’s no good reason to take just Lefty to be extra or just Righty to be extra.

Perhaps we should choose as if each of Lefty’s and Righty’s welfare is worth *half* the equivalent amount of non-extra welfare. That seems like a natural move, and it could be justified by appeal to answer (4): each of Lefty and Righty is “half-extra.” The move could also be justified by appeal to answer (5): one of Lefty and Righty is extra, but it is indeterminate which.²³ We need then only add a couple more claims: (a) faced with this kind of indeterminacy, we should choose as if there’s a 0.5 probability that it’s Lefty that’s extra and a 0.5 probability that it’s Righty, and (b) we should be risk-neutral with respect to these probabilities.

However, choosing as if Lefty’s and Righty’s welfare is worth half the equivalent in non-extra welfare is also hard to square with the Person-Affecting Restriction. To see why, consider *Benign A-Fission*:²⁴

Benign A-Fission

No Split		Split B		Split C	
Anna	80	Anna	10	Anna	10
Lefty	Ω	Lefty	90	Lefty	60
Righty	Ω	Righty	10	Righty	60

Split B seems better for Anna than No Split, for the reasons given above. In particular, Split B would be better for Anna than No Split if only Lefty existed, and it’s difficult to see how reintroducing Righty could make Anna worse off: Pre-fission Anna shouldn’t think that she’d benefit by bribing the surgeon to drop the right half of her brain, thereby ensuring that Righty doesn’t exist (Campbell ms: 9; Nozick 1981: 64–65). After all, the relation that matters is plausibly *intrinsic* (Parfit 1984: 263): Whether Lefty’s welfare contributes to the quality of Anna’s life—and the degree to which it does so—depends only on the relations that obtain between Lefty and Anna. It doesn’t depend on

23. Johansson (2010) suggests this view about personal identity in fission cases: one of Lefty and Righty is identical to Anna, but it is indeterminate which.

24. The coming argument draws on Huemer’s (2008: 901–903) Benign Addition Argument, inspired by Parfit’s (1984: ch. 19) original Mere Addition Paradox. It shares some similarities with Gustafsson’s and Kosonen’s (2024) argument that, all else equal, it is better for a person to have more fission-products with good lives.

what happens elsewhere, or on the relations that obtain between either Lefty or Anna and any other person.²⁵

Split C, meanwhile, seems better for Anna than Split B. Lefty's life-branch is a little worse in Split C, but Righty's life-branch is much better. Split C is more equal, and it has greater total and average welfare. Given the transitivity of 'better for,' the result is that Split C is better for Anna than No Split.

Suppose that, nevertheless, we choose No Split over Split C. If we should choose as if each of Lefty's and Righty's welfare is worth half the equivalent in non-extra welfare, hard and very hard views imply that we were required to make that choice, while soft views imply that we were at least permitted to do so. These person-affecting views thus seem undeserving of the name, since Split C is better for the only non-extra person and very good for everyone else.²⁶

That leaves only answer (6): Person-affecting views should designate neither Lefty nor Righty as extra. This answer avoids any impersonal-seeming implications. The catch is that (6) exposes person-affecting views to analogues of *all* of the problems faced by impersonal views like total utilitarianism. Take

25. To pump intuitions here, suppose that Anna has the opportunity to incur some welfare-cost in her own life-episode in order to bestow some welfare-benefit on Lefty's life-episode. Per my intuitions, Anna need not determine whether Righty will exist in order to determine whether the cost is prudentially worth paying.

One might doubt this claim if one thinks that Anna should have some fixed degree of prudential concern for all her splittees, so that Lefty's welfare is prudentially half as important to Anna if Righty is on the scene (and vice versa). But this view has counterintuitive implications, especially in cases of negative welfare. Here's an example. Suppose that—by default—only Lefty will exist, with welfare -100. Anna can ensure that Righty exists too, but doing so will drop Lefty's welfare to -180, and Righty will have welfare -10. That change seems worse for Anna, but the view under consideration implies that it's a change for the better.

26. One might think that advocates of *multiple occupancy* can avoid this conclusion. On the multiple occupancy interpretation of fission cases, both splittees exist prior to fission as distinct, co-located persons (Lewis 1976). One might then suggest that Righty lives a life of welfare 20 if we choose Split B and lives a life of welfare 80 if we choose No Split. Since Righty would then be worse off in Split B than in No Split, a requirement to choose No Split over Split B would not fall foul of the Person-Affecting Restriction.

The first thing to note is that this suggestion departs from the orthodox multiple occupancy view. On the orthodox view, Righty doesn't exist if we choose No Split, and so isn't worse off if we choose Split B. One could adopt a revised multiple occupancy view on which each of Lefty and Righty exist even in No Split, but this view delivers implausible verdicts in other cases. Suppose for example that in No Split, Anna's welfare score is -100 (and hence, on this revised multiple occupancy view, Lefty's and Righty's welfare scores are also -100). Suppose also that in Split D, Anna is split immediately, and Lefty's and Righty's welfare scores are -99. Given the revised multiple occupancy view's interpretation of the case, Split D is better than No Split for both Lefty and Righty, and so any plausible moral view will require us to choose Split D. But on a more natural understanding of the case, choosing Split D means nearly doubling the suffering that occurs, for no gain whatsoever. That gives us reason to reject the revised multiple occupancy view.

See Briggs and Nolan (2015) for other implausible consequences stemming from multiple occupancy views.

Repugnant Transition, for example. Total utilitarianism requires that we make the transition, while person-affecting views do not.²⁷ But now consider a minor variation, which we can call *Repugnant Fission*.²⁸ Suppose that the world contains only people at the start of their lives. Suppose also that we have two options. We can leave these people unsplit, in which case their lives will be wonderful. Alternatively, we can immediately split each of these people many times, in which case each splittee's life-branch would be barely worth living. If each splittee is non-extra, their welfare counts in the usual non-extra way. Extant person-affecting views typically aggregate non-extra welfare by summing: a population *A* is at least as good as a population *B* with respect to non-extra welfare iff *A* contains at least as great a sum-total of non-extra welfare as *B* (Cohen 2020; Horton 2021: 499; McDermott 1982; 2019; Meacham 2012; Thomas 2023). Thus if the number of splittee is great enough (and if we're required to do what's best), these person-affecting views imply that we're required to choose fission.²⁹ But this verdict seems about as repugnant as Repugnant Transition. After all, the post-fission world could be almost exactly like the post-transition world. Both could contain a vast number of human beings subsisting on "muzak and potatoes" (Parfit 1986: 148).

More generally, wherever *creating new people* raises a problem for impersonal views, *creating new splittees* raises an analogous problem for person-affecting views coupled with (6): the claim that splittees are non-extra. For example, while person-affecting views are largely neutral about making happy people, (6) implies that they're in favour of making happy splittees. All else equal, creating

27. Strong actualism is (something of) an exception. See fn 12.

28. It's an analogue of Parfit's (1984: 388) Repugnant Conclusion. I discuss a similar point in (Thornley 2025). There are two other analogues of the Repugnant Conclusion worth mentioning. The first—Nebel's (2019a)—concerns the probability that a person exists. The second—discussed by McTaggart (1927: 452–453), Crisp (2006: 112) and Temkin (2009: §7; 2012: 119) among others—concerns the duration that a person exists. This latter analogue is counterintuitive, and person-affecting views and impersonal views are similarly vulnerable to it. Whether it's implied depends entirely on how the relevant view aggregates momentary welfare into lifetime welfare. Because person-affecting views can imply this analogue, one might think that avoiding Repugnant Transition isn't such a significant advantage of person-affecting views: although person-affecting views avoid Repugnant Transition itself, problems of the same kind can remain. I think an anonymous referee for this point.

29. One might claim that each splittee is identical to all the other splittees along with the original person from whom they split (Dainton 1992), and that the welfare of life-branches is intrapersonally-aggregated in such a way that each original person is worse off in the post-fission population no matter how many splittees they spawn. But identity-relations this pervasive lead to all kinds of trouble. Setting aside familiar implications about the possibility of one person in two bodies unwittingly playing tennis against herself (Parfit 1984: 256–257), the ethical upshots also seem tough-to-swallow. One might have to agree that harming Lefty to benefit Righty is no more morally fraught than harming Anna on Monday to benefit Anna on Tuesday. And even if fission preserves identity, repeated iterations of the Benign A-Fission Argument can be used to conclude that, given enough splittees, each of the original people is better off in the post-fission population.

happy splittees is required. These person-affecting views thus contravene an analogue of Narveson's Slogan. Person-affecting views paired with (6) also violate an analogue of Roberts' and Chappell's Deeper Intuition. According to *Deeper Intu-Fission*, we ought to benefit an existing person in some fission-free way by some amount g rather than create a new splittee with welfare g . But if splittees are non-extra and we aggregate non-extra welfare by summing, creating the new splittee is permissible. And if the new splittee would have welfare ever-so-slightly-greater-than- g , creating them would be required. Like impersonal views, then, person-affecting views paired with (6) imply something that might seem perverse: in some circumstances, we could be required to create new splittees rather than help those suffering today.

Answer (6)—the claim that splittees are non-extra—also means that person-affecting views can require large sacrifices in the present for the sake of unlikely benefits in the far future. Consider again the Good Future from §3, in which Earth plays host to 10^{17} wonderful lives. This time, however, imagine that humans reproduce a little more like amoebae. We split after 70 years, with one splittee dying soon afterwards and the other living 70 years before themselves splitting, and so on. Suppose that there are 10^{17} 70-year life-branches in this population and 10^{17} fleeting life-branches. Each 70-year life-branch is wonderful, and each fleeting life-branch is neutral. Each splittee is fully-connected-in-all-the-ways-that-might-matter to the person from whom they split. The Short Future also features humans-like-amoebae but is otherwise as before: runaway climate change, nuclear war, the release of an engineered pathogen, or some other disaster causes humanity to go extinct a century from now.

Suppose again that we—the present generation—can shift the probabilities with which these two futures come about. By all worsening our (by-default wonderful) current life-branches so that they're just barely worth living, we can decrease the chance of the Short Future by p and increase the chance of the Good Future by p . The other option is business-as-usual. For what values of p must we take the plunge? If our person-affecting view take splittees to be non-extra, aggregates non-extra welfare by summing, ranks risky options using expected value theory, and requires us to do what's best, the answer is roughly "Any value greater than or equal to 0.0000008." We're required to make enormous sacrifices in the present-day for the sake of far-future splittees that may never exist, even if those sacrifices have just an eight-in-ten-million chance of paying off. Call this implication *Fission Sacrifice*. It seems to me about as implausible as Our Sacrifice: expected total utilitarianism's verdict in our original case.

Here's the current state-of-play. If—as I've claimed—Repugnant Fission is about as implausible as Repugnant Transition, violations of Deeper Intu-Fission are about as implausible as violations of the Deeper Intuition, and Fission Sacrifice is about as implausible as Our Sacrifice, then the advantages of a certain

family of person-affecting views over expected total utilitarianism have evaporated. This family of person-affecting views consists of those views that embrace (6)—the claim that splittees are non-extra—along with aggregation-by-summing, expected value theory, and a requirement to do what’s best. We have little reason to prefer these person-affecting views to expected total utilitarianism, and we have more-than-little reason for the opposite preference. Person-affecting views face problems that impersonal views do not.

One might reply that person-affecting views and answer (6) are blameless in these cases: The real culprit(s) in Repugnant Fission and the violation of Deeper Intu-Fission is aggregation-by-summing or maximising consequentialism (the claim that we’re required to do what’s best), and the real culprit(s) in Fission Sacrifice are aggregation-by-summing, expected value theory, or maximising consequentialism. This thought has some merit: There are alternative aggregation rules, rules for ranking risky options, and rules for deriving requirements that allow person-affecting views paired with (6) to avoid some of these problems. However, this fact is cold comfort for advocates of person-affecting views, because those very same aggregation rules (albeit applied to all welfare, rather than just non-extra welfare), rules for ranking risky options, and rules for deriving requirements allow impersonal views to avoid those same problems, along with their non-fission analogues. And in fact, answer (6) implies a more general conclusion: no matter what aggregation rule *A*, rule for ranking risky options *R*, and rule for deriving requirements *D* we choose, each person-affecting view paired with (6), *A*, *R*, and *D* will face fission analogues of whatever problems exist for an impersonal view paired with *A*, *R*, and *D*. If these fission analogues are as implausible as the originals, we have little reason to prefer person-affecting views plus (6) over the corresponding impersonal views.

One might then claim that the fission analogues are more plausible than the originals. One might defend this claim by pointing out that Fission Sacrifice isn’t really a sacrifice, at least not in any moral sense. That’s because we—the present generation—are connected-in-all-the-ways-that-might-matter to these far-future splittees. Their existence would be good *for us*, and to such an extent that we’re each better off in expectation choosing fission sacrifice over business-as-usual. One might say something similar about Repugnant Fission: splitting is better *for each person in the original population*. Although each of the resulting life-branches is barely worth living, each of the original people is connected-in-all-the-ways-that-might-matter to many such life-branches. That makes splitting better for them overall. One might also claim that violating Deeper Intu-Fission isn’t so bad, because creating a new splittee with welfare *g* is *a way of benefitting the original person by g*.

With regards to the relative plausibility of Fission Sacrifice and Our Sacrifice, I’ve run out of arguments. I can only report my own view, which is that

the appeal to *betterness for us* doesn't make much difference. The lion's share of implausibility—in both cases—comes from the enormous upfront cost and the tiny probability of any payoff. Faced with this pricey long-shot bet, I get little solace from the thought that it will be *I*—rather than someone else—who might get to enjoy wonderful life-branches far into the future.

What about the relative plausibility of Deeper Intu-Fission and the Deeper Intuition? Here I have an argument. Although it's true that creating a new splittee with welfare *g* is a way of benefitting the original person by *g*, this point does little to address the most troubling violations of Deeper Intu-Fission: cases in which we ought to create a new happy splittee from an existing person rather than relieve the suffering of a *different* existing person. The possibility of these cases seems to me at least as implausible as the possibility of cases that violate the Deeper Intuition: cases in which we ought to create a new happy person rather than relieve the suffering of an existing person.

Finally, consider the relative plausibility of Repugnant Fission and Repugnant Transition. Here too I have an argument. Although splitting is better for each person in the original population, we need also consider the interests of each splittee. And doing so illuminates a sense in which Repugnant Fission is *less* plausible than Repugnant Transition. As noted above, many of the people living mediocre lives in Repugnant Transition are plucked from the ether. They're not connected-in-any-way-that-might-matter to any person who would have existed if we made the other choice. But all of the splittees living mediocre life-branches in Repugnant Fission are connected-in-all-the-ways-that-might-matter to a person who would have lived a wonderful life if we made the other choice.

The upshot, I claim, is that the implications of person-affecting views paired with (6), aggregation-by-summing, expected value theory, and maximising consequentialism remain counterintuitive, and roughly *as* counterintuitive as the corresponding implications of expected total utilitarianism. It's counterintuitive to suppose that a population of people living wonderful lives is worse than a large population of humans on life-branches barely worth living, *even if those humans are the product of fission*. It's counterintuitive to suppose that we must sacrifice all that's good in this century for the sake of a long-shot bet on far-future welfare, *even in a world where we reproduce like amoebae*. And it's counterintuitive to suppose that we should create new humans rather than help those suffering today, *even if these new humans are split off from those already existing*. If all that's the case, then it can't be the impersonal aspect of expected total utilitarianism that's cause for concern. The trouble—if there is any—must have its roots elsewhere: in expected total utilitarianism's demandingness, its indifference to injustice, its happy substitution of quality for quantity, its taste for speculative gambles, or its inhuman patience. And any person-affecting view

worth the name—paired with expected total utilitarianism’s rules for aggregation, ranking risky options, and deriving requirements—*also* has these features. The upshot is that we have little reason to prefer these person-affecting views to expected total utilitarianism.

What’s more, as distasteful as the above features may seem, we know that there are costs to denying aggregation-by-summing, expected value theory, and maximising consequentialism.³⁰ Person-affecting views paired with answer (6) are liable to bear these costs, just as impersonal views are. Therefore, whatever rules we settle on for aggregating welfare, ranking risky options and deriving requirements, we have little reason to pair these rules with a person-affecting view rather than an impersonal view. If the former view is truly person-affecting, it will face fission analogues of all of the problems that afflict the impersonal view, in addition to the problems that afflict person-affecting views alone.

5. Conclusion

On first glance, person-affecting views seem to have significant advantages over impersonal views like expected total utilitarianism. Person-affecting views tend to respect Narveson’s Slogan and the Deeper Intuition, and they mostly avoid implications like Repugnant Transition and Our Sacrifice.

In this paper, I’ve argued that the advantages of person-affecting views aren’t as significant as they first appear. Parfit-style fission cases confront these views with a dilemma: either violate the spirit of the Person-Affecting Requirement, or else imply fission analogues of the problems that blight impersonal views. In light of this dilemma, the advantages of person-affecting views largely evaporate. What remains are the problems unique to them.

Acknowledgements

For helpful comments and discussion, I thank Jonas H. Aaron, Adam Bales, Tomi Francis, Hilary Graves, Teruji Thomas, Hayden Wilkinson, and two anonymous referees for *Ergo*.

30. I’ve been emphasising the problems for expected total utilitarianism, but various impossibility theorems prove that every aggregation rule (Arrhenius 2011; Carlson 1998; Parfit 1984: ch. 19) and rule for ranking risky options (Beckstead & Thomas 2024) has at least one implausible-seeming implication (see also Thornley 2021). Departures from maximising consequentialism face issues like the paradox of deontology (Scheffler 1982; 1985), various paradoxes of supererogation (Horton 2017; Kamm 1985; 2001: ch. 12; 2007: 30–31; Muñoz 2020; Sidgwick 1907: 220), and Kagan’s (1989) objections to constraints and prerogatives.

References

- Arrhenius, Gustaf (2011). The Impossibility of a Satisfactory Population Ethics. In Ehtibar N. Dzhanfarov & Lacey Perry (Eds.), *Descriptive and Normative Approaches to Human Behavior* (1–26). World Scientific Publishing Company.
- Arrhenius, Gustaf (forthcoming). *Population Ethics: The Challenge of Future Generations*. Oxford University Press.
- Bader, Ralf M. (2022). Person-Affecting Utilitarianism. In Gustaf Arrhenius, Krister Bykvist, Tim Campbell, & Elizabeth Finneron-Burns (Eds.), *The Oxford Handbook of Population Ethics* (251–270). Oxford University Press.
- Beckstead, Nick (2013). *On the Overwhelming Importance of Shaping the Far Future* (PhD Thesis, Rutgers University). Retrieved from <http://dx.doi.org/doi:10.7282/T35M649T>
- Beckstead, Nick and Teruji Thomas (2024). A Paradox for Tiny Probabilities and Enormous Values. *Noûs*, 58(2), 431–455.
- Bigelow, John and Robert Pargetter (1988). Morality, Potential Persons and Abortion. *American Philosophical Quarterly*, 25(2), 173–181.
- Bradley, Ben (2013). Asymmetries in Benefiting, Harming and Creating. *The Journal of Ethics*, 17(1), 37–49.
- Briggs, Rachael and Daniel Nolan (2015). Utility Monsters for the Fission Age. *Pacific Philosophical Quarterly*, 96(3), 392–407.
- Broome, John (1999). *Ethics Out of Economics*. Cambridge University Press.
- Campbell, Tim (ms). *Personal Identity and Aggregation*. Retrieved from http://www.academia.edu/8854258/Personal_Identity_and_Aggregation
- Carlson, Erik (1998). Mere Addition and Two Trilemmas of Population Ethics. *Economics & Philosophy*, 14(2), 283–306.
- Chappell, Richard Yetter (2015). Value Receptacles. *Noûs*, 49(2), 322–332.
- Chappell, Richard Yetter (2017). Rethinking the Asymmetry. *Canadian Journal of Philosophy*, 47(2–3), 167–177.
- Cohen, Daniel (2020). An Actualist Explanation of the Procreation Asymmetry. *Utilitas*, 32(1), 70–89.
- Crisp, Roger (2006). *Reasons and the Good*. Oxford University Press.
- Dainton, Barry F. (1992). Time and division. *Ratio*, 5(2), 102–128. <https://doi.org/10.1111/j.1467-9329.1992.tb00041.x>
- Francis, Tomi (2021). Anonymity and Non-Identity Cases. *Analysis*, 81(4), 632–639. <https://doi.org/10.1093/analys/anab031>
- Francis, Tomi (ms). *In Favour of Making Happy People*. Retrieved from <https://drive.google.com/file/d/1XoVvHokd6st5oVspvLzZVZdA6uVKbk38/view>
- Frick, Johann (2017). On the Survival of Humanity. *Canadian Journal of Philosophy*, 47(2–3), 344–367.
- Frick, Johann (2020). Conditional Reasons and the Procreation Asymmetry. *Philosophical Perspectives*, 34(1), 53–87.
- Greaves, Hilary (2017). Population Axiology. *Philosophy Compass*, 12(11).
- Greaves, Hilary and John Cusbert (2022). Comparing Existence and Non-Existence. In Jeff McMahan, Tim Campbell, James Goodrich, & Ketan Ramakrishnan (Eds.), *Ethics and Existence: The Legacy of Derek Parfit* (61–92). Oxford University Press.
- Greaves, Hilary and William MacAskill (2025). The Case for Strong Longtermism. In Hilary Greaves, Jacob Barrett, and David Thorstad (Eds.), *Essays on Longtermism* (17–49). Oxford University Press.

- Gustafsson, Johan E. and Petra Kosonen (2024). Do Lefty and Righty Matter More Than Lefty Alone? *Erkenntnis*, 89, 1921–1926. <https://link.springer.com/article/10.1007/s10670-022-00610-1>
- Hare, Caspar (2007). Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist? *Ethics*, 117(3), 498–523.
- Heyd, David (1988). Procreation and value: Can Ethics Deal with Futurity Problems? *Philosophia*, 18(2–3), 151–170.
- Heyd, David (1992). *Genethics: Moral Issues in the Creation of People*. University of California Press.
- Holtug, Nils (2004). Person-Affecting Moralities. In Torbjörn Tännsjö & Jesper Ryberg (Eds.), *The Repugnant Conclusion: Essays on Population Ethics* (129–161). Kluwer Academic Publishers.
- Holtug, Nils (2010). *Persons, Interests, and Justice*. Oxford University Press. Retrieved from <https://academic.oup.com/book/9142>
- Horton, Joe (2017). The All or Nothing Problem. *The Journal of Philosophy*, 114(2), 94–104.
- Horton, Joe (2021). New and Improvable Lives. *The Journal of Philosophy*, 118(9), 486–503.
- Huemer, Michael (2008). In Defence of Repugnance. *Mind*, 117(468), 899–933. <https://doi.org/10.1093/mind/fzn079>
- Hume, David (1738/2007). *A Treatise of Human Nature* (David Fate Norton and Mary J. Norton, Eds.). Oxford University Press.
- Johansson, Jens (2010). Parfit on Fission. *Philosophical Studies*, 150(1), 21–35.
- Kagan, Shelly (1989). *The Limits of Morality*. Clarendon Press.
- Kamm, Frances Myrna (1985). Supererogation and Obligation. *The Journal of Philosophy*, 82(3), 118–138.
- Kamm, Frances Myrna (2001). *Morality, Mortality: Volume II: Rights, Duties, and Status*. Oxford University Press.
- Kamm, Frances Myrna (2005). Moral Status and Personal Identity: Clones, Embryos, and Future Generations. *Social Philosophy and Policy*, 22(2), 283–307.
- Kamm, Frances Myrna (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford University Press.
- Lewis, David (1976). Survival and Identity. In Amelie Oksenberg Rorty (Ed.), *The Identities of Persons* (17–40). University of California Press.
- McDermott, Michael (1982). Utility and population. *Philosophical Studies*, 42(2), 163–177.
- McDermott, Michael (2019). Harms and objections. *Analysis*, 79(3), 436–448.
- McTaggart, John M.E. (1927/1969). *The Nature of Existence Volume II*. (Ed. C. D. Broad). Cambridge University Press.
- Meacham, Christopher J. G. (2012). Person-Affecting Views and Saturating Counterpart Relations. *Philosophical Studies*, 158(2), 257–287.
- Mogensen, Andreas (2019). Staking Our Future: Deontic Long-Termism and the Non-Identity Problem. *GPI Working Paper, No. 9-2019*. Retrieved from https://www.global-prioritiesinstitute.org/wp-content/uploads/2019/Mogensen_Staking_Our_Future.pdf
- Muñoz, Daniel (2020). Three paradoxes of supererogation. *Noûs*, 55(3), 1–18. <https://doi.org/10.1111/nous.12326>
- Narveson, Jan (1973). Moral problems of population. *The Monist*, 57(1), 62–86.
- Nebel, Jacob M. (2019a). An Intrapersonal Addition Paradox. *Ethics*, 129(2), 309–343.
- Nebel, Jacob M. (2019b). Asymmetries in the Value of Existence. *Philosophical Perspectives*, 33(1), 126–145.

- Nebel, Jacob M. (2021). Conservatism about the Valuable. *Australasian Journal of Philosophy*, 1–15.
- Nozick, Robert (1981). *Philosophical Explanations*. Harvard University Press.
- Ord, Toby (2020). *The Precipice: Existential Risk and the Future of Humanity*. Bloomsbury.
- Parfit, Derek (1984). *Reasons and Persons*. Clarendon Press.
- Parfit, Derek (1986). Overpopulation and the Quality of Life. In Peter Singer (Ed.), *Applied Ethics* (145–164). Oxford University Press.
- Parsons, Josh (2002). Axiological Actualism. *Australasian Journal of Philosophy*, 80(2), 137–147.
- Podgorski, Abelard (2023). Complaints and Tournament Population Ethics. *Philosophy and Phenomenological Research*, 106(2), 344–367. <https://doi.org/10.1111/phpr.12860>
- Pummer, Theron (2024). Future Suffering and the Non-Identity Problem. *GPI Working Paper, No. 17-2024*. Retrieved from <https://www.globalprioritiesinstitute.org/wp-content/uploads/Pummer-T-Future-Suffering-and-the-Non-Identity-Problem.pdf>
- Roberts, Melinda A. (2011a). An Asymmetry in the Ethics of Procreation. *Philosophy Compass*, 6(11), 765–776.
- Roberts, Melinda A. (2011b). The Asymmetry: A Solution. *Theoria*, 77(4), 333–367.
- Ross, Jacob (2015). Rethinking the Person-Affecting Principle. *Journal of Moral Philosophy*, 12(4), 428–461.
- Scheffler, Samuel (1982). *The Rejection of Consequentialism*. Clarendon Press.
- Scheffler, Samuel (1985). Agent-Centred Restrictions, Rationality, and the Virtues. *Mind*, XCIV(375), 409–419.
- Setiya, Kieran (2014). The Ethics of Existence. *Philosophical Perspectives*, 28(1), 291–301.
- Sidgwick, Henry (1907/1981). *The Methods of Ethics* (Seventh Edition). Macmillan.
- Singer, Peter (2011). *Practical Ethics* (5th ed.). Cambridge University Press.
- Spencer, Jack (2021). The Procreative Asymmetry and the Impossibility of Elusive Permission. *Philosophical Studies*, 178(11), 3819–3842.
- Temkin, Larry S. (2009). Aggregation Within Lives. *Social Philosophy and Policy*, 26(1), 1–29.
- Temkin, Larry S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.
- Thomas, Teruji (2016). *Topics in Population Ethics*. (Doctoral dissertation, University of Oxford). Retrieved from <https://philpapers.org/archive/THOTIP.pdf>
- Thomas, Teruji (2023). The Asymmetry, Uncertainty, and the Long Term. *Philosophy and Phenomenological Research*, 107(2), 470–500.
- Thornley, Elliott (2021). The Impossibility of a Satisfactory Population Prospect Axiology (Independently of Finite Fine-Grainedness). *Philosophical Studies*, 178, 3671–3695. <https://link.springer.com/article/10.1007/s11098-021-01621-4>
- Thornley, Elliott (2023). The Procreation Asymmetry, Improvable-Life Avoidance, and Impairable-Life Acceptance. *Analysis*, 83(3), 517–526.
- Thornley, Elliott (2024). A Non-Identity Dilemma for Person-Affecting Views. *GPI Working Paper, No. 6-2024*. Retrieved from <https://www.globalprioritiesinstitute.org/wp-content/uploads/Elliott-Thornley-A-non-identity-dilemma-for-person-affecting-views.pdf>
- Thornley, Elliott (2025). Critical-Set Views, Biographical Identity, and the Long Term. *Australasian Journal of Philosophy*, 103(3), 638–653. <https://www.tandfonline.com/doi/full/10.1080/00048402.2025.2476692>
- Warren, Mary Anne (1977). Do Potential People Have Moral Rights? *Canadian Journal of Philosophy*, 7(2), 275–289.