



The procreative asymmetry and the impossibility of elusive permission

Jack Spencer¹

Accepted: 3 March 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract This paper develops a form of moral actualism that can explain the procreative asymmetry. Along the way, it defends and explains the attractive asymmetry: the claim that although an impermissible option can be self-conditionally permissible, a permissible option cannot be self-conditionally impermissible.

Keywords Procreative asymmetry · Actualism · Normative variance · Ethics · Epistemology · Deontology

1 Introduction

There is an apparent asymmetry in the moral status of choices not to procreate. On the one hand, we do not seem to violate a moral obligation when we choose not to create a sure-to-be-happy person. Consider:

Joy or Nothing. We are deciding whether to create Joy. We know that Joy, if created, would lead a happy life, and that nobody other than Joy will be affected by our choice.

Morality seems to permit what we've done if we do not create Joy—and the permission does not seem to be owed to the usual trappings of procreation. We may suppose that we could have created Joy just by pressing a button, and that Joy, if created, would never have causally interacted with any of us. Even with these additional suppositions, the following seems true:

✉ Jack Spencer
jackspen@mit.edu

¹ Department of Phil. and Ling., MIT, 32-d929, 77 Mass. Ave, Cambridge, MA 02139, USA

- (1) Not creating Joy is permissible if we do not create Joy.

On the other hand, we seem to satisfy a moral obligation when we choose not to create a sure-to-be-miserable person. Consider:

Misery or Nothing. We are deciding whether to create Misery. We know that Misery, if created, would lead a miserable life, and that nobody other than Misery will be affected by our choice.

The following seems true:

- (2) Not creating Misery is obligatory if we do not create Misery.

I take the “procreative asymmetry” to be the conjunction of (1) and (2),¹ and one can see straightaway why it’s puzzling.² It’s as if morality cares about the misery that would have filled Misery’s life but is indifferent to the joy that would have filled Joy’s. One goal of this paper is to develop a way of explaining the procreative asymmetry.

My proposed explanation amends an existing proposal, a view which, following Hare (2007), I call “strong actualism”. Strong actualism—which I’ll characterize more precisely in the next section—is motivated by the *person-affecting intuition*: the claim that nothing can be better or worse unless it’s better or worse for actual people. As we’ll see, strong actualism can explain the “happy half” of the procreative asymmetry, i.e., (1). But it has a fatal structural defect,³ admits of clear counterexamples,⁴ and cannot explain the “miserable half” of the procreative asymmetry, i.e., (2).⁵ My proposal, which I call “stable actualism”, is better. It rectifies the structural defect, avoids the clear counterexamples, and can explain both halves of the procreative asymmetry. It will not appeal to everyone. Those who does not find the person-affecting intuition compelling are unlikely to accept it. But

¹ One could think that there is also an asymmetry in the moral status of choices to procreate, accepting not just (1) and (2), but also:

- (3) Not creating Joy is permissible if we create Joy.
 (4) Not creating Misery is obligatory if we create Misery.

One who accepted (1), (2), (3), and (4) might take the procreative asymmetry to be the conjunction of those four claims. But, as I say in Sect. 3, I reject (3).

The procreative asymmetry is sometimes stated in terms of reasons; see *e.g.* Chappell (2017: 167), Frick (2020), and McMahan (1981: 100; 2009: 49). Like Wedgwood (2015), I’m somewhat skeptical of “reasons”-talk, so I state it in deontic terms.

² For other work on the procreative asymmetry, see *e.g.* Algander (2015), Boonin (2014), Broome (2004; 2005), Bykvist (2007a; 2007b), Chappell (2017), Elstein (2005), Frick (2020), Hare (2007; 2011), Harman (2004; 2009), Heyd (1992), Holtug (2001), McMahan (1981; 1994; 2009), Narveson (1967; 1973; 1978); Parfit (1982; 1987), Persson (2009), Roberts (2003a; 2010; 2011a; 2011b; 2019), and Sterba (1987).

³ *Cf.* Bykvist (2007b) and Hare (2007).

⁴ For criticisms of strong actualism, see *e.g.* Arrhenius (2009; 2015), Bykvist (2007b), Frick (2020), Hare (2007), McMahan (1981; 1994), Roberts (2010: ch. 2; 2011a; 2011b), and Parfit (1982; 1987).

⁵ *Cf.* McMahan (1981: 102) and Roberts (2011a: Sect. 6; 2011b: 771).

the modest claims I make on behalf of stable actualism—that it can explain both halves of the procreative asymmetry, and that it's superior to strong actualism—can be accepted even by those who do not find the person-affecting intuition compelling.

The route from strong actualism to stable actualism is paved by a much less modest thesis I seek to defend concerning the nature of permission and the scope of normative variance.

There is normative variance when permissibility depends on choice.⁶ For example, if $A = \{a_1, \dots, a_n\}$ is set of options available to the agent, and $C(a_i)$ is the set of options that are permissible if option a_i is chosen, then the choice the agent faces involves normative variance just if, for some $a_i, a_j \in A$, $C(a_i) \neq C(a_j)$.

One putative variety of normative variance is the sort had by options that support their own permissibility. If $a_i \in C(a_i)$, but $a_i \notin C(a_j)$, for some $a_j \in A$, we'll say that a_i is *attractively permissible*. An attractively permissible option is permissible if chosen. Another putative variety of normative variance is the sort had by options that oppose their own permissibility. If $a_i \notin C(a_i)$, but $a_i \in C(a_j)$, for some $a_j \in A$, we'll say that a_i is *elusively permissible*. An elusively permissible option is impermissible if chosen.

Strong actualism give rise to both varieties. In *Joy or Nothing*, it predicts that not creating Joy is attractively permissible: that it's permissible to not create Joy if and only if we do not create Joy. In *Misery or Nothing*, it predicts that creating Misery is elusively permissible: that it's permissible to create Misery if and only if we do not create Misery.

I accept the possibility of attractive permission. The possibility of attractive permission is a crucial part of what allows strong actualism to explain the happy half of the procreative asymmetry, and I think that strong actualism correctly explains the happy half of the procreative asymmetry. But I reject the possibility of elusive permission. If $a_i \notin C(a_i)$, then, for any $a_j \in A$, $a_i \notin C(a_j)$. The less modest thesis I seek to defend may be called the *attractive asymmetry*—the claim that, although attractive permission is possible, elusive permission is not.⁷

Stable actualism can be thought of as the view that results from subtracting the possibility of elusive permission from strong actualism. It's a view for those who find the person-affecting intuition compelling, but are convinced, as I am, that nothing can be elusively permissible.

⁶ See Carlson (1995: ch 6), who credits both the idea and the name 'normative variance' to Wlodek Rabinowicz. Also see Prichard (1968: 37) and Bykvist (2007b: 99).

⁷ Bykvist (2007b) lists many normative theories that are committed to the possibility of both attractive and elusive permission. As we'll see in Sect. 7, the claim that an agent is permitted to believe any proposition that is likely on her evidence entails the possibility of both attractive and elusive permission. Philosophers who reject the possibility of elusive permission usually also reject the possibility of attractive permission; see e.g. Broome (2004: 74), Carlson (1995: ch. 6), Frick (2020), Hare (2007; 2011: n. 11), and Roberts (2010: 62).

2 Strong actualism

Hare (2007) draws a distinction between strong actualism and what he calls “weak actualism”.⁸ We’ll look at weak actualism in Sect. 6, but I want to begin with strong actualism and some of its merits and demerits.

Strong actualism is a perspectival moral theory, which allows the value of an option to depend on which option is chosen. Let $A = \{a_1, \dots, a_n\}$ be set of the options available to the agent,⁹ and, to make things simple, let’s suppose that, for each $a \in A$, there is some uniquely closest a -world, which is the world that would be actual were the agent to choose a . Co-indexing in the natural way, we then have the set of *actualizable* worlds, $W = \{w_1, \dots, w_n\}$. If we let $Va_j(a_i)$ be the a_j -value of option a_i —that is, the value of option a_i , if option a_j were chosen—then we can visualize strong actualism by constructing the following $n \times n$ matrix:

$$\begin{array}{cccc} & a_1 & \cdots & a_n \\ a_1 & Va_1(a_1) & \cdots & Va_1(a_n) \\ \cdots & \cdots & \cdots & \cdots \\ a_n & Va_n(a_1) & \cdots & Va_n(a_n) \end{array}$$

When the value of an option depends on which option is chosen, the entries in the column that represent the option will vary across the rows.

Strong actualism, as I’ll be understanding it here, has three main tenets. The first is a qualified Pareto principle, which is meant to capture the person-affecting intuition:¹⁰

Perspectival Pareto. Let a_i and a_j be any two options in A , and let a_k be any option in A . Then:

- i. If a_i would have been better than a_j for someone and would not have been worse than a_j for anyone, if a_k were chosen, then, things being appropriately equal, $Va_k(a_i)$ is greater than $Va_k(a_j)$.
- ii. If a_i would have been worse than a_j for someone and would not have been better than a_j for anyone, if a_k were chosen, then, things being appropriately equal, $Va_k(a_i)$ is less than $Va_k(a_j)$.
- iii. If a_i would have been equally as good as a_j for everyone if a_k were chosen, then, things being appropriately equal, $Va_k(a_i)$ is equal to $Va_k(a_j)$.

The qualification, “things being appropriately equal”, is added because there is (or anyway very well might be) more to morality than individuals’ interests. There may

⁸ Bykvist (2006: 275–6) calls weak actualism, “ratificationism”.

⁹ I assume that which options the agent has never depends on which option the agent chooses. Also, to ensure that agents always choose exactly one option, I assume that options are mutually exclusive and jointly exhaustive.

¹⁰ This Pareto principle resembles the principle Hare (2007: 502) calls “Minimal Commitment”.

be rights, or matters of desert, or global values, like equality and relative priority.¹¹ But with regard to the questions in procreative ethics that I want to explore here, I am going to assume that things are appropriately equal.¹² (Those who think otherwise may read this essay as an attempt to explain the procreative asymmetry just by appeal to individuals' interests.)

The second tenet of strong actualism is a principle about the value of existence, which I'll call,

Comparability. Necessarily, for any person x , a possible world in which x leads a happy life is better for x than a possible world in which x does not exist, and a possible world in which x does not exist is better for x than a possible world in which x leads a miserable life.¹³

Comparability enjoys considerable intuitive support, but it's controversial because it conflicts with two *prima facie* plausible claims. Let x be some actual miserable person. Let $w_{@}$ be the actual world, and let w be some possible world in which x does not exist. According to Comparability, w is better for x than $w_{@}$. But the following claim is tempting:

Accessibility. If w is better for x than $w_{@}$, then, if w had been actual, w (still) would have been better for x than $w_{@}$.¹⁴

And nothing can be better or worse for something that does not exist. So,

Not Counterfactually Better. It is not the case that, if w had been actual, w (still) would have been better for x than $w_{@}$.

Accessibility and Not Counterfactually Better entail, *contra* Comparability, that w is not better for x than $w_{@}$.

Like Holtug (2001) and Arrhenius & Rabinowicz (2010), I reject Accessibility. An analogy with preference might be helpful. An actual person can prefer a world in which they never exist to the actual world: $w_{@}$ may satisfy x 's (actual) preferences less so than does w , even though there is no such thing as x 's preferences at w because x does not exist at w . And the same holds true for what is good for someone: $w_{@}$ may realize x 's (actual) interests less so than does w , even though there is no such thing as x 's interests at w because x does not exist at w .¹⁵

¹¹ The conflict between the unqualified version of Perspectival Pareto and egalitarian intuitions is brought out by Arrhenius (2015), who also shows that a weakening of Perspectival Pareto, an inequality aversion principle, and an egalitarian dominance principle together entail the repugnant conclusion. One theory that gives pride of place to relative priority is Buchak (2017).

¹² In making this assumption, I follow Hare (2007).

¹³ If people are necessary beings, as necessitists, like Williamson (2013), maintain, then we should replace talk of existence with talk of chunkiness.

¹⁴ The name for this principle comes from Bykvist (2007a: 348). Its proponents include Broome (1999: 168), Bykvist (2007a), McMahan (2009), and Parfit (1987: 489). Its opponents include Adler (2009), Arrhenius and Rabinowicz (2010), Holtug (2001), Roberts (1998; 2003a; 2011a; 2011b), and Shiffrin (1999).

¹⁵ Comparability also conflicts with the following two claims: (1) if w is better for x than $w_{@}$, then x 's welfare in w exceeds x 's welfare in $w_{@}$; and (2) if x does not exist at w , then it is not the case that x 's

The third tenet of strong actualism, a principle that might seem to go without saying, connects value to permissibility:

Maximization. The options that are permissible at world w are all and only the options that maximize value at w .

Options that maximize value at the actual world will be said to maximize value, *sans phrase*. Thus, according to Maximization, an option is (actually) permissible if and only if it maximizes value (*sans phrase*).¹⁶

3 Joy

One of the merits of strong actualism is that it can explain the happy half of the procreative asymmetry, i.e., (1). According to strong actualism, *Joy or Nothing* exhibits normative variance. If we create Joy, then creating Joy is obligatory, and if we do not create Joy, then both options are permissible. Underlying these perspectival deontic facts are perspectival values. Creating Joy is self-conditionally obligatory (i.e. obligatory if chosen) because, if we create Joy, then creating Joy is better for someone (viz. Joy) and not worse for anyone. Not creating Joy is self-conditionally permissible (i.e. permissible if chosen) because, if we do not create Joy, then not creating Joy is equally as good as creating Joy for everyone.¹⁷

Footnote 15 continued

welfare in w exceeds x 's welfare in $w_{@}$. But, as Johansson (2010) and Arrhenius and Rabinowicz (2010) argue, (1) is false. How good w is for x is determined, not by x 's welfare in w or by the degree to which x 's interests in w are satisfied, but by the degree to which w (actually) satisfies x 's (actual) interests. And as I say in the text, $w_{@}$ may satisfy x 's (actual) interests less so than does w , even though there is no such thing as x 's interests at w because x does not exist at w .

The failure of Accessibility brings contingency in its wake. The two comparative claims—that possible worlds in which x leads a happy life are better for x than possible worlds in which x does not exist, and that possible worlds in which x does not exist are better for x than possible worlds in which x leads a miserable life—are only contingently true, since they are true only at worlds at which x exists. This contingency does not falsify Comparability. But it does falsify the following stronger principle:

Necessary Comparability. Necessarily, for any person x , necessarily, a possible world in which x leads a happy life is better for x than a possible world in which x does not exist, and a possible world in which x does not exist is better for x than a possible world in which x leads a miserable life.

Thanks to a helpful referee for pressing me on this point.

¹⁶ Hare (2007: 503), who uses 'S@' to refer to actual people, characterizes strong actualism as follows: "*Strong Actualism*.—The moral status of any a_j , actual or not, is determined by whether its outcome is better or worse for people in S@ than the outcomes of the other available actions).

¹⁷ Of course one option being obligatory excludes another being permissible. But if normative variance is possible, then one option being self-conditionally obligatory does not exclude another being self-conditionally permissible; for the fact that one option is the only permissible option if chosen does not entail that another option cannot be permissible if chosen. In *Joy or Nothing*, for example, if a_J and $a_{\bar{J}}$ are the options of creating and not creating Joy, respectively, then it could be the case that $C(a_J) = \{a_J\}$ and $C(a_{\bar{J}}) = \{a_J, a_{\bar{J}}\}$.

It's important that the procreative asymmetry here concerns the moral status of choices not to procreate. One could think that there is also an asymmetry in the moral status of choices to procreate, accepting, in addition to (1) and (2), both of the following:

- (3) Not creating Joy is permissible if we create Joy.
- (4) Not creating Misery is obligatory if we create Misery.

Someone who accepts (1), (2), (3), and (4) may be inclined to identify the procreative asymmetry with the conjunction of these four claims, and thus would identify the happy half of the procreative asymmetry with the conjunction of (1) and (3). Strong actualism then would not be capable of explaining the happy half of the procreative asymmetry; for strong actualism, though consistent with (1), (2), and (4), is inconsistent with (3).

But it's a mistake to identify the happy half of the procreative asymmetry with the conjunction of (1) and (3). The intuition that constitutes the happy half of the procreative asymmetry is a felt absence of transgression—the intuition that nobody acts impermissibly by not creating Joy. The conjunction of (1) and (3) entails that nobody acts impermissibly by not creating Joy, but (3) does no work: (1) entails, all by itself, that nobody acts impermissibly by not creating Joy. What (3) adds to (1) is an extraneous claim about the deontic status of unchosen options: (1) says that not creating Joy is permissible whenever chosen; (3) says that not creating Joy is permissible whenever unchosen. And not only is (3) extraneous to the intuition that constitutes the happy half of the procreative asymmetry, it's false. Not creating Joy, though permissible if we do not create Joy, is impermissible if we create Joy; for, if we create Joy, then creating Joy is better for someone (*viz.* Joy) and not worse for anyone.

It's tempting to think that the procreative asymmetry says that we're never obligated to create Joy, but that's not quite right. What the procreative asymmetry says is that the obligation to create Joy is never violated. And, as strong actualism rightly says, the reason why the obligation to create Joy is never violated is that we are obligated to create Joy only if we create Joy.

The fact that we're obligated to Joy if we create Joy makes itself known, not just in relations of Pareto optimality, but also in retrospection, I think. There is an interesting retrospective asymmetry in *Joy or Nothing*. If we do not create Joy, then both options retrospectively appear choiceworthy; the fitting retrospective attitude is neither regret, nor gladness, but something more akin to retrospective ambivalence. If we create Joy, however, then the only option that retrospectively appears choiceworthy is creating Joy; the fitting retrospective attitude is gladness. There are, of course, many ways one could try to explain this retrospective asymmetry, but the explanation that strong actualism offers is appealingly simple. According to strong actualism, the retrospective asymmetry is a reflection of the normative variance. The options that retrospectively appear choiceworthy are exactly the permissible options.

My inclination to accept this connection between retrospective choiceworthiness and permissibility reinforces my belief that strong actualism handles *Joy or Nothing* correctly.

4 Misery

Strong actualism does not handle *Misery or Nothing* correctly, however. According to strong actualism, *Misery or Nothing* exhibits normative variance. If we create Misery, then we're obligated to not create Misery, and if we do not create Misery, then both options are permissible. In making these predictions, strong actualism errs twice. It wrongly predicts that *Misery or Nothing* exhibits normative variance, and it wrongly predicts that we're permitted to create Misery if we don't.

The retrospective asymmetry in *Joy or Nothing* is notably absent in *Misery or Nothing*. Unlike what strong actualism predicts, creating Misery retrospectively appears unchoiceworthy, no matter which option we choose.

We thus have a counterexample to strong actualism. And once we've seen one, it's easy to construct others. Consider:

Misery or Moremisery. We are deciding whether to create Misery or another person, Moremisery. We know that Misery, if created, would lead a miserable life; that Moremisery, if created, would lead an even more miserable life; and that nobody other than Misery or Moremisery will be affected by our choice.

What to say about *Misery or Moremisery* is not entirely clear to me. I'm open to the view that it's a dilemma, in which both options are unconditionally impermissible, and I'm open to the view that creating Misery is unconditionally obligatory. But I'm not open to what strong actualism says. According to strong actualism, *Misery or Moremisery* exhibits normative variance. We're obligated to create whoever we do not create. But the claim that we're obligated to create Moremisery if we create Misery is, I think, clearly false.¹⁸

5 Deontic consistency

Not only does strong actualism admit of counterexamples; it also violates a deontic consistency principle that should be affirmed.

There is a *dilemma* if none of the agent's options are permissible. The most familiar deontic consistency principle is:

No Dilemmas. Dilemmas are impossible.

¹⁸ Roberts (2010: ch.2) argues that two other examples make trouble for strong actualism.

The first is the "Basic Case", in which option a_1 makes A exist with welfare 100 and B not exist and option a_2 makes A exist with welfare 0 and B exist with welfare 100. My intuitions differ from Roberts about this case. I think that strong actualism is right: that a_1 is obligatory if chosen, and that both options are permissible if a_2 is chosen.

The second is "Addition Plus," in which option a_1 makes A exist with welfare 10 and B not exist, option a_2 makes A exist with welfare 11 and B exist with welfare 1, and option a_3 makes A and B both exist with welfare 5. Addition Plus turns partly on considerations of equality, which I've set aside. But it's worth pointing out that if a_3 is better both from the perspective of the world that would be actual if a_2 were chosen and from the perspective of the world that would be actual if a_3 were chosen, then the view I call "hierarchical actualism" delivers the result that a_3 is obligatory, no matter which option is chosen.

No Dilemmas is controversial. Its opponents often point to Sophie's choice scenarios, in which an agent must decide which of two innocent people will suffer some horrible fate.¹⁹ An even stronger challenge to No Dilemmas comes from symmetric miserable creation cases, however. Consider:

Misery or Equalmisery. We're deciding whether to create Misery or another person, Equalmisery. We know that Misery, if created, would lead a miserable life; that Equalmisery, if created, would lead an equally miserable life; and that nobody other than Misery or Equalmisery will be affected by our choice.

In Sophie's choice scenarios, the chosen option is better for *someone*: namely, the one who does not suffer the horrible fate. In *Misery or Equalmisery*, the chosen option is better for *no-one*, and much, much worse for the only person it affects. A strong case can be made that *Misery or Equalmisery* is a dilemma.

No Dilemmas is not, however, at this point in the dialectic, the relevant principle. Strong actualism is, in fact, consistent with No Dilemmas.

There is a *weak dilemma* if some of the agent's options are permissible, but none of the agent's options are self-conditionally permissible. What's disturbing about strong actualism is that it conflicts with:

No Weak Dilemmas. Weak dilemmas are impossible.

The conflict is brought out by *Misery or Equalmisery*. According to strong actualism, if we create Misery, the only permissible option is creating Equalmisery, and if we create Equalmisery, the only permissible option is creating Misery. Some option is permissible, but no option is self-conditionally permissible.

Weak dilemmas are, to put it mildly, odd. Here's Bykvist bringing out the oddity:

What is especially troublesome is [...] a case where, if you did A, A would be wrong and not doing A right, whereas if you did not do A, A would be right and not doing A wrong. For this situation involves unavoidable wrong-doing in the sense that whatever you were to do, you would do something that would be wrong. You are damned if you do, damned if you don't. Or more exactly, you *would* be damned if you *were* to do it, and you *would* be damned if you *were* not to do it. [...]. Normally, a dilemma is seen as a situation in which all available actions *are* wrong. This is not the situation here. No matter how you act, there is an available act that *is* right. If A is performed, then refraining from doing A is right; if A is not performed, A is right. But this is cold comfort. For you cannot act in such a way that were you to act in that way you would comply with the theory. (Bykvist, 2007b: 116–7).

Here's Hare bringing out the oddity:

[If you face a weak dilemma, then you are] in the odd position of knowing, in advance of having made up your mind about what to do, that the action you will take is the one you ought not to take, and the action you could take but won't is the one you ought to take. You are *weakly fated* to do what you ought

¹⁹ See McConnell (2018) and references therein.

not to do. It's not that you can't avoid doing what you ought not to do; it's just that you know that you actually won't. (Hare, 2007: 507).

Both Bykvist and Hare regard the conflict between strong actualism and No Weak Dilemmas as a reason to reject strong actualism, and I think they're right. No Weak Dilemmas should be affirmed.

6 Stable Actualism

Strong actualism admits of counterexamples, so at least one of its three tenets are false. I'm inclined to accept both Perspectival Pareto and Comparability, so I'm inclined to reject Maximization.²⁰

One principle that we could replace Maximization with leads to the view that Hare calls "weak actualism". Recall that the a_j -value of a_i , $Va_j(a_i)$, is the value of option a_i if option a_j were chosen. Let's say that option a_i maximizes *self-conditional value* just if the a_i -value of a_i is not exceeded by the a_i -value of any other option; that is, an option maximizes self-conditional value just if it would maximize value if chosen. Instead of Maximization, weak actualists accept:

Self-Conditional Maximization. The permissible options are all and only the options that maximize self-conditional value.²¹

One merit of weak actualism is that it handles *Misery or Nothing* correctly. Since not creating Misery maximizes self-conditional value, and since creating Misery does not, weak actualism rightly predicts that we are unconditionally obligated to not create Misery.²²

A second merit of weak actualism is that it's consistent with No Weak Dilemmas. Whether an option maximizes self-conditional value never depends on which option is chosen, so weak actualism is consistent with the impossibility of normative variance. The impossibility of normative variance entails the impossibility of elusive permission, and the impossibility of elusive permission entails No Weak Dilemmas.

²⁰ Some might be tempted at this juncture to motivate a distinction between "positive" and "negative" value, and then try to argue that the person-affecting intuition holds only of positive value; cf. Parfit (1987: 525–6). On the resultant view: In *Joy or Nothing*, if we do not create Joy, then not creating Joy is not better because it's not better for anyone; but in *Misery or Nothing*, if we do not create Misery, then creating Misery is worse even though it's not worse for anyone. The alleged distinction between "positive" and "negative" value is mysterious, however, as is the claim that the person-affecting restriction applies to one but not the other. I think that a better option is to retain Perspectival Pareto, embrace the claim that creating Misery is not worse if we do not create Misery (since, after all, it's then not worse for anyone), and respond to the counterexamples by rejecting Maximization.

²¹ Hare (2007: 502–3), who uses ' Sa_j ' to refer to the people who are would be actual if a_j were chosen, characterizes weak actualism as follows: "*Weak Actualism.*—The moral status of any a_j , actual or not, is determined by whether its outcome is better or worse for people in Sa_j than the outcomes of the other available actions."

²² Another merit of weak actualism: It avoids the prediction, in *Misery or Moremisery*, that we are obligated to create Moremisery if we create Misery.

A third merit of weak actualism is that it can explain both halves of the procreative asymmetry. Weak actualism entails (2) because, in *Misery or Nothing*, the only option that maximizes self-conditional value is not creating Misery. Weak actualism entails (1) because, in *Joy or Nothing*, both options maximize self-conditional value.

Despite these merits, I think we should reject weak actualism, however; for I think that weak actualism handles *Joy or Nothing* incorrectly. Weak actualism predicts that not creating Joy is *unconditionally* permissible, and that prediction is mistaken. Not creating Joy is permissible if, but only if, Joy is not created.

Weak actualists are right, I think, to draw attention to the maximization of self-conditional value, but wrong, I think, about its import. Whether an option maximizes self-conditional value is relevant, not because options are made permissible by maximizing self-conditional value, but because an option cannot *stably* maximize value without maximizing self-conditional value.

Let $M(a_j)$ be the options that maximize a_j -value. Let a_w be the option chosen at world w , and let a_{a} be the option that is actually chosen, whichever it is. Option a_i maximizes value at world w just if it's a member of $M(a_w)$, and *stably* maximizes value at w just if it's a member of both $M(a_w)$ and $M(a_i)$. Option a_i maximizes value (*sans phrase*) just if it's a member of $M(a_{\text{a}})$, and *stably* maximizes value (*sans phrase*) just if it's a member of both $M(a_{\text{a}})$ and $M(a_i)$. An option that does not stably maximize value will be said to *elusively* maximize value, since, although the option maximizes value, the agent would not have maximized value had they chosen it.

Instead of Maximization, I think we should accept some conception of permissibility that entails:

Stable Maximization. If some option stably maximizes value at w , then the options that are permissible at w are all and only the options that stably maximize value at w .

Stable actualism is the view whose tenets are Perspectival Pareto, Comparability, and Stable Maximization. (One naturally wonders how to complete stable actualism, where a completion would conjoin to Stable Maximization a specification of what it takes for an option to be permissible when no option stably maximizes value. In Sect. 10, I consider two possible completions. But, for now, to remain neutral among the various possible completions, I want to focus on stable actualism itself, incomplete though it is.)

Like strong actualism, stable actualism handles *Joy or Nothing* correctly. Stable actualism and strong actualism coincide when all of the options that maximize value also stably maximize value, and, in *Joy or Nothing*, all of the options that maximize value also stably maximize value. If we create Joy, the only option that stably maximizes value is creating Joy.²³ If we do not create Joy, then both options stably maximize value.²⁴ Stable actualism thus rightly predicts that not

²³ If we create Joy, $M(a_{\text{a}}) = M(a_J)$, and $a_J \in M(a_J)$ and $a_{\bar{J}} \notin M(a_J)$.

²⁴ If we do not create Joy, $M(a_{\text{a}}) = M(a_{\bar{J}})$, and $a_J \in M(a_{\bar{J}})$, $a_{\bar{J}} \in M(a_{\bar{J}})$, and $a_J \in M(a_J)$.

creating Joy is self-conditionally permissible and that creating Joy is self-conditionally obligatory.

Stable actualism also has all of the aforementioned merits of weak actualism. It handles *Misery or Nothing* correctly, since, no matter which option is chosen, the only option that stably maximizes value is not creating Misery.²⁵ It's consistent with No Weak Dilemmas, since it's consistent with the impossibility of elusive permission. And it can explain both halves of the proactive asymmetry—it entails both (1) and (2).

These merits give us reason to take stable actualism seriously. Over the next four sections, I'll consider four objections to it. The first alleges that stable actualism is ad hoc. The second alleges that stable actualism, like any theory that gives rise to normative variance, should be rejected. The third alleges that stable actualism is refuted by certain nonidentity cases. And the fourth alleges that there is no plausible way to complete stable actualism.

7 The impossibility of elusive permission

Isn't stable actualism ad hoc? What principled reason could there be for rejecting Maximization in favor of Stable Maximization?

Answer: We should reject Maximization in favor of Stable Maximization because we should accept the following principle:

Stability. If p makes it the case that an agent is permitted to choose a , then p would have obtained had the agent chosen a .²⁶

Stability imposes a possible use condition on permission-making. It says that a fact cannot make an agent permitted to choose a unless it can make the agent permitted to have chosen a . Stability does not impose a possible *non-use* condition on permission-making. The fact that makes an agent permitted to choose a need not hold at all of the actualizable worlds; indeed, it could hold at only one of the actualizable worlds. But it must hold at the world that the agent would actualize by choosing a , since otherwise the agent could not *use* the permission it makes.

Stability is compatible with the claim that the maximization of a quantity is both necessary and sufficient for permissibility if the quantity in question is stably maximized whenever it's maximized. But if a quantity can be maximized without being stably maximized, then Stability is inconsistent with the maximization of the quantity being both necessary and sufficient for permissibility. Maximization and Stable Maximization are both concerned with the quantity that we've been calling "value", and that quantity, as we've seen, can be maximized without being stably

²⁵ If we create Misery, then $M(a_{\text{misery}}) = M(a_M)$, and $a_{\overline{M}} \in M(a_M)$ and $a_M \notin M(a_M)$. If we do not create Misery, then $M(a_{\text{no misery}}) = M(a_{\overline{M}})$, and $a_M \in M(a_{\overline{M}})$, $a_{\overline{M}} \in M(a_{\overline{M}})$, and $a_M \notin M(a_M)$.

²⁶ Stability resembles a principle that Hare (2011: 196) calls, "Reasons are not Self-Undermining". But the principles are not equivalent—in fact, I reject Hare's principle. Stability is also connected to the conception of guidance developed in Spencer and Wells (2019) and Spencer (forthcoming).

maximized.²⁷ Stability thus provides a principled reason for favoring Stable Maximization over Maximization.

I believe that Stability holds, not just of moral permission, but of every kind of permission. Consider epistemic permission, for example.

There is a close connection between high evidential probability and epistemic permission. Usually, if p is likely on an agent's evidence, the agent is permitted to believe that p . This connection between high evidential probability and epistemic permission holds even when a proposition has high evidential probability only because the agent believes it. Consider a case of confirmed reliability. The agent has an impermissible belief. Although p is unlikely on her evidence, she believes that p . Moreover, she knows that she believes that p . She then gains some new evidence, which makes it very likely that p is true if and only if she believes that p . The new evidence makes it permissible for her to believe that p , even though she would not be permitted to believe that p if she did not believe that p .

But there are cases in which a proposition is likely on an agent's evidence only because the agent does not believe the proposition, and in those cases the high evidential probability of the proposition does not seem to make it permissible for the agent to believe the proposition. Let me offer two illustrations.

First, a case of confirmed unreliability. The agent has a permissible belief. She believes that $\sim p$, and $\sim p$ is likely on her evidence. Moreover, she knows that she believes that $\sim p$. She then gains some new evidence, which makes it very likely that p is true if and only if she believes that $\sim p$. As a result, p is likely on her evidence. But there is a strong intuition that she is not permitted to believe that p , despite the fact that p is likely on her evidence.

Second, a case of disconfirming belief. The evidential probability of a proposition might be high, even though the evidential probability of the proposition conditional on the agent believing the proposition is low. Adapting an example Kotzen (MS) uses in discussion of desire-as-belief, let p be the proposition that I would be a good politician. The evidential probability of p might be high—perhaps I've been plain-dealing, heretofore. But the evidential probability of p conditional on me believing that p might be low, since believing that I would be a good politician might be strong evidence that I would not be. If the evidential probability of p is high, but the evidential probability of p conditional on me believing that p is low, then it seems that I am not permitted to believe that p , despite the high probability that p enjoys on my evidence.

What underlies these epistemic intuitions, I think, is the analog of Stability. We think that a fact cannot make it permissible for an agent to believe a proposition if the fact would not obtain if the agent believed the proposition. (If each proposition is an "option" and an agent "chooses" an option by believing it, then Stability

²⁷ Of course, it's controversial whether the quantity that we have been calling "value" is relevant to the permissibility of options. Weak actualists, for example, maintain that the quantity that is relevant to the permissibility of options is not value, but rather self-conditional value. But, in *Joy or Nothing*, not creating Joy maximizes value if and only if we do not create Joy, so those who think, as I do, that not creating Joy is permissible if and only if we do not create Joy have reason to think that the maximization of value is indeed relevant to the permissibility of options.

covers both the moral and the epistemic cases.) A proposition has high evidential probability *stably* if it has high evidential probability both in the actual world and in the world that would be actual if the agent believed the proposition. It may be the case that an agent is permitted to believe any proposition that has high evidential probability stably. But the mere fact that a proposition has high evidential probability does not entail that the agent is permitted to believe the proposition, for the proposition might have high evidential probability *elusively*, as the case of confirmed unreliability and the case of disconfirming belief illustrate.²⁸

If Stability is true, then, with the help of two ancillary principles, we can derive the impossibility of elusive permission. Here's the first ancillary principle:

Necessitation. If p makes it the case that the agent is permitted to choose a , then p necessitates that the agent is permitted to choose a .

If we think of permission-makers as the grounds of permissions, then Necessitation is just an instance of the widely accepted claim that grounds necessitate what they ground.²⁹

The second principle requires a new bit of terminology. If an agent is permitted to choose a and there is some fact, p , that makes it the case that the agent is permitted to choose a , then we'll say that the agent's permission to choose a is *derivative*. Here's, the second ancillary principle:

Derivative Elusive Permissions. Every elusive permission is derivative.

Stability and Necessitation together entail that no derivative permission is elusive. If an agent is permitted to choose a , and p makes it the case that the agent is permitted to choose a , then, by Stability, p would have obtained had the agent chosen a . So, by Necessitation, the agent's permission is not elusive—the agent would have been permitted to choose a had the agent chosen a . Thus, if Derivative Elusive Permissions holds, it follows that elusive permission is impossible.

I think all permissions are derivative, so I think Derivative Elusive Permissions follows from a more general principle. But even if I countenanced primitive permissions, I would not countenance primitive elusive or attractive permissions. Elusive and attractive permissions depend in a special way on the agent's choice, but there would be no way to explain this special dependence if the permissions were primitive. So not only do I accept Derivative Elusive Permissions, I also accept:

Derivative Attractive Permissions. Every attractive permission is derivative.

²⁸ A similar phenomenon is familiar in decision theory. It's tempting to think that any option that maximizes (causal) expected value is rationally permissible for an agent to choose. But there are cases—like Egan's (2007) *Psychopath Button* and Ahmed's (2014) *Dicing with Death*—in which an option maximizes expected value only because the agent is confident that she will not choose the option, and in such cases the option that maximizes expected value does not seem to be rationally permissible. A number of authors have responded to these cases by defending a stability condition, arguing that an option is made rationally permissible by maximizing expected value only if the option also maximizes expected value conditional on its being chosen; see *e.g.* Egan (2007), Harper (1986), Spencer and Wells (2019), and Spencer (forthcoming).

²⁹ See *e.g.* Rosen (2010: 118).

But while these four principles—Stability, Necessitation, Derivative Elusive Permissions, and Derivative Attractive Permissions—entail that elusive permission is impossible, they do not entail that attractive permission is impossible. Stability is asymmetric. It imposes a possible use condition on permission-making, but it does not impose a possible non-use condition. The asymmetry of Stability thus paves the way for the attractive asymmetry.

Stability also lays to rest the first objection to stable actualism. The move from Maximization to Stable Maximization is not ad hoc.

8 The possibility of normative variance

Like strong actualism, stable actualism entails that normative variance is possible. Some philosophers think that normative variance is not possible³⁰:

It is quite implausible that what one ought to do depends on what one does. I think this is enough to cast severe doubt on actualism. (Broome, 2004: 74).

But why think that normative variance is impossible? What's implausible about it?³¹

In a particularly incisive discussion, Bykvist (2007b) identifies two things that are potentially problematic about theories that give rise to normative variance.

The first we've seen already. Bykvist thinks that any normative theory must satisfy both No Dilemmas and No Weak Dilemmas, and he points out that many normative theories that give rise to normative variance conflict with these principles. I'm skeptical of No Dilemmas, as I said. But I think Bykvist is right that we should reject any normative theory that conflicts with No Weak Dilemmas, and I go one small step further: I think we should reject any normative theory that predicts the possibility of elusive permission. (Any theory that conflicts with No Weak Dilemmas predicts the possibility of elusive permissions, but the reverse is not true.)

As Bykvist points out, however, there is no argumentative route from these deontic consistency principles to the impossibility of normative variance. Indeed, in Sect. 10, I'll offer one completion of stable actualism that's consistent with both No Dilemmas and No Weak Dilemmas.

The other problem Bykvist discusses concerns deliberation. Bykvist, who uses "NI" to abbreviate the thesis that normative variance is impossible, puts the point as follows:

[A] theory that violates NI is a poor *guide* to action. One might take this to be a decisive argument for NI for the following reason. When you use a theory as a guide to action, you use the theory in your deliberations about what to do. On the basis of this deliberation you then make up your mind and decide what to do. But if your theory violates NI, then in order to decide whether an action

³⁰ Also see e.g. Hare (1975: 219) and Narveson (1978: 44).

³¹ Both Bykvist (2007b) and Howard-Snyder (2008) defend the possibility of normative variance.

has a certain normative status [...] you have to know whether or not you are going to perform it. But there is no point in deliberating about whether to perform an action if either you believe that you will perform it, or you believe that you will not perform it. If you believe that you will perform the action, the issue is settled for you, and there is no point in deliberating about it further. If you believe that you will not perform the action, the action is no longer a serious possibility, i.e., something that is compatible with what you believe [...] so again there is no point in deliberating about whether to perform it. (Bykvist, 2007b: 110–1).³²

Normative variance inhibits deliberation—Bykvist is right about that. But is that a good reason to reject the possibility of normative variance? I think it isn't.

The point of deliberation is to avoid impermissible options, and the possibility of normative variance is an impediment to that goal. In a case of normative variance, a deliberating agent cannot, simply by deliberating, winnow her options down to just the permissible ones. But the agent *can*, simply by deliberating, winnow her options down to just the self-conditionally permissible ones. And since elusive permission is impossible, permissibility entails self-conditional permissibility. In a case of normative variance, then, a deliberating agent *can*, simply by deliberating, winnow her options down to a superset of the permissible options, where every member of that set is permissible if chosen.

The question, then, is whether that's enough. And to my mind, it clearly is. If it weren't enough, then a case like *Joy or Nothing* would be problematic. In *Joy or Nothing*, if we create Joy, then not creating Joy is impermissible, but we cannot, simply by deliberating, eliminate the option of not creating Joy. But this consequence does not seem problematic; rather, it seems exactly right. There are normative differences between *Joy or Nothing* and the choice between bales of hay that Buridan's ass faces, but, *vis-à-vis* deliberation, I think the two should pattern together. It takes an act of will, over and above sound deliberation, for Buridan's ass to choose a bale of hay, and I think that, similarly, in *Joy or Nothing*, if we've not decided whether we will create Joy, it takes an act of will, over and above sound deliberation, to choose an option.

I'm not aware of other arguments against the possibility of normative variance,³³ but a number of people who have seen this paper have wondered at this point about the relationship between normative variance and practical reasoning.³⁴ It's natural to think that there is some connection between reasons and reasoning: that *p* can be a reason for an agent to choose *a* only if there is some sound bit of practical reasoning that could lead the agent from *p* to the choosing of *a*.³⁵ If permission-

³² Here Bykvist echoes Carlson (1995: 101).

³³ For some additional arguments in favor of the possibility of normative variance, see Howard-Snyder (2008).

³⁴ Thanks to Caspar Hare, Daniel Muñoz, and Kieran Setiya for discussion.

³⁵ See *e.g.* Hare (2011), Setiya (2014a), and Williams (1981).

makers are reasons, and this connection between reasons and reasoning holds,³⁶ then we have the following constraint on permission-making:

Practical Reasoning. If p makes an agent permitted to choose a , then there is some sound bit of practical reasoning that could lead the agent from p to the choosing of a .

One might wonder whether Practical Reasoning is consistent with the possibility of normative variance.

Practical Reasoning is not consistent with the possibility of elusive permission.³⁷ A sound bit of practical reasoning has no false premises. So, if p does not hold at the world that the agent would actualize by choosing a , then there is no sound bit of practical reasoning that could lead the agent from p to the choosing of a . Indeed, Practical Reasoning might help to explain why elusive permission is impossible. But so far as I can tell, Practical Reasoning is perfectly consistent with the possibility of attractive permission. In *Joy or Nothing*, the fact that not creating Joy maximizes value could be our reason for not creating Joy, and that seems to entail that there is a sound bit of practical reasoning that could lead us from the fact that not creating Joy maximizes value to not creating Joy.

Having looked at all of the arguments against the possibility of normative variance of which I am aware, and having found all of them wanting, I conclude that there is no sound argument against the possibility of normative variance.

I also conclude that there really is an attractive asymmetry. Already we have seen two potential illustrations of attractive permission: *Joy or Nothing* and the case of confirmed reliability. These cases convince me that attractive permission is possible. But let me briefly mention three other potential illustrations.

First, abortion. Harman (1999) argues that a fetus has moral status if and only if the fetus actually develops into a person. If Harman is right, then abortion illustrates attractive permission—aborting a fetus is permissible only if the fetus is aborted.

Second, adoption. Suppose that an agent can adopt A or B; that the agent would improve A's life to a considerable degree by adopting A; and that the agent would improve B's life to a slightly greater degree by adopting B. If it's permissible for an agent to give a smaller benefit to their own child instead of a larger benefit to a child that is not their own, then suboptimal adoption illustrates attractive permission—adopting A is permissible only if A is adopted.

Third, prudence. An agent is deciding whether to move to Town or Glad City. If the agent moves to Glad City, she will form new, pro-city preferences, and will thus prefer Glad City to Town. If the agent moves to Town, she will form new preferences that do not tell between towns and cities, and will thus be indifferent between Glad City and Town. Say that an option maximizes *actual desirability* if no option satisfies the agent's actual preferences to a greater degree. It's tempting to think that an option is prudentially permissible if it (stably) maximizes actual desirability. But if an option is prudentially permissible if it (stably) maximizes

³⁶ It's not obvious that permission-makers are reasons; thanks to Kieran Setiya for discussion.

³⁷ Cf. Hare (2011: 196; 2011: n. 11).

actual desirability, then the choice between Town and Glad City has the same normative structure as *Joy or Nothing*. Moving to Glad City is unconditionally permissible, and moving to Town is permissible only if the agent moves to Town.

None of these potential illustrations is irresistible. But if there is no *general* argument against the possibility of attractive permission, then, given the diversity of potential illustrations, someone who wants to resist the possibility of attractive permission has their work cut out for them.

9 Lesserjoy

Some opponents of strong actualism think it delivers the wrong results in nonidentity cases like the following³⁸:

Joy or Lesserjoy. We are deciding whether to create Joy or another person, Lesserjoy. We know that Joy, if created, would lead a happy life; that Lesserjoy, if created, would lead a happy, but less happy, life, on account of intermittent misery. And we know that nobody other than Joy or Lesserjoy will be affected by our choice.

According to strong actualism, *Joy or Lesserjoy* exhibits normative variance. We're obligated to create whoever we create. But some philosophers defend the *improvement claim*: that we are *unconditionally* obligated to create Joy.³⁹ If the improvement claim is true, then strong actualism is false.

The improvement claim is also inconsistent with stable actualism. In *Joy or Lesserjoy*, the chosen option, whichever it is, is the only option that stably maximizes value. So, like strong actualism, stable actualism predicts that we're obligated to create whoever we create.

The fact that strong and stable actualism stand and fall together on this point should not surprise. Both are motivated by the person-affecting intuition, and rejecting the improvement claim is just part of the cost one pays in affirming the person-affecting intuition. In *Joy or Lesserjoy*, the chosen option, whichever it is, is better than the unchosen option for someone and not worse for anyone.

The modest claims that I make on behalf of stable actualism—that it can explain both halves of the procreative asymmetry, and that it's superior to strong actualism—are not threatened by the improvement claim. But, that said, I think the improvement claim is more doubtful than it's usually presumed to be.

One concern comes from retrospection. If the improvement claim is true, then we could know, prior to creating Lesserjoy, that creating Lesserjoy is impermissible. It's natural to think that retrospective gladness and permissibility co-vary: that we should not be glad to have chosen an option that we knew to be impermissible prior

³⁸ The literature on the nonidentity problem, and the conflict between it and the person-affecting intuition, is vast. See Roberts (2019) and the references therein.

³⁹ There are other nonidentity cases that enjoy greater intuitive support. But those cases—like the depletion case, in Parfit (1987: 312), and the slave child case, in Kavka (1982: 101)—involve considerations that are excluded by the “things being appropriately equal” qualification.

to choosing.⁴⁰ But in order to accommodate the retrospective asymmetry in *Joy or Lesserjoy*, proponents of the improvement claim must deny that retrospective gladness and permissibility co-vary. If we create Joy, the fitting retrospective attitude is gladness. But the same is true in reverse. If we create Lesserjoy, we should not wish that we had instead created Joy; the fitting retrospective attitude is gladness, not retrospective ambivalence or regret.⁴¹

So, if the improvement claim is true, then there will be cases in which it is appropriate to be glad to have chosen an option that we knew to be impermissible prior to choosing—and that’s odd.

The pattern of appropriate retrospection in *Joy or Lesserjoy* exactly parallels the normative variance that strong and stable actualism predict, so strong and stable actualism avoid the oddity. It’s not surprising, by the lights of strong and stable actualism, that gladness is appropriate if we have created Lesserjoy; for gladness is the appropriate retrospective attitude to have when an obligatory option has been chosen, and, according to strong and stable actualism, creating Lesserjoy is self-conditionally obligatory.

Another concern comes from the tension between the improvement claim and the procreative asymmetry.⁴²

We can bring the tension out informally by dividing *Joy or Lesserjoy*. Suppose that instead of making one decision, we make two. We decide whether to create Joy, and we decide whether to create Lesserjoy. The procreative asymmetry entails that both options in both choices are self-conditionally permissible. But the claim that both options in both choices are self-conditionally permissible is hard to square with the improvement claim. If it’s permissible to not create Joy (in *Joy or Nothing*), and it’s permissible to create Lesserjoy (in *Lesserjoy or Nothing*), then it’s puzzling why it should be impermissible to not create Joy and create Lesserjoy (in *Joy or Lesserjoy*).

⁴⁰ For arguments that retrospective gladness and permissibility do not co-vary, see Harman (2009) and Setiya (2014b).

⁴¹ In *Joy or Lesserjoy*, I claim that, if we have created Lesserjoy, then retrospective gladness is appropriate. There is an attenuated sort of gladness that is appropriate whenever something has a positive component. For example, if I have a choice between \$5 and \$10, and I choose \$5, then, although regret is appropriate, attenuated gladness might also be appropriate, since I am, after all, \$5 richer on account of the choice I’ve made. But the sort of gladness that I am concerned—what might be called “all-things-considered” gladness—is not compatible with regret. It’s not appropriate to be (all-things-considered) glad to have chosen \$5, but I claim that it *is* appropriate to be (all-things-considered) glad to have created Lesserjoy.

One could deny that retrospective (all-things-considered) gladness is appropriate if we have created Lesserjoy. I find that claim intuitively compelling, but I have not argued for it here. It may be worth noting, however, that there are others who also find that claim compelling. For example, although Setiya (2014b) and I disagree on several of points, we both agree that, if Lesserjoy has been created, then retrospective (all-things-considered) gladness is appropriate.

⁴² For other discussions of the interaction between the improvement claim and the procreative asymmetry, see e.g. Boonin (2014: ch. 7), Frick (2020), Harman (2004), Heyd (1992), McMahan (1981), Narveson (1973; 1978), Roberts (2010; 2011a; 2019), and Parfit (1987).

One way to formalize this tension appeals to two principles.⁴³ Let A and B be two sets of options. If some $a \in A$ is compossible with each $b \in B$ and self-conditionally permissible relative to A , no matter which $b \in B$ is chosen, then we'll say that a is, *independently of B , self-conditionally permissible relative to A* . We then have the first principle⁴⁴:

Weak Independent Agglomeration. If some $a \in A$ is, independently of B , self-conditionally permissible relative to A , and some $b \in B$ is, independently of A , self-conditionally permissible relative to B , then $a \& b$ is self-conditionally permissible relative to $A \times B$.⁴⁵

The second principle is a variation of Sen's (1970) Alpha:

Self-Conditional Alpha. If a is one of the self-conditionally permissible options in A , and A^- is a subset of A that contains a , then a is one of the self-conditionally permissible options in A^- .

Neither principle is apodictic,⁴⁶ but both are plausible. And given the two principles, the improvement claim is inconsistent with the procreative asymmetry. Let $J = \{a_J, a_{\bar{J}}\}$ be the set that contains the options of creating and not creating Joy, and let $L = \{a_L, a_{\bar{L}}\}$ be the set that contains the options of creating and not creating Lesserjoy. Each option in each set is compossible with each option in the other set. The procreative asymmetry entails that each option in J is, independently of L , self-conditionally permissible relative to J , and that each option in L is, independently of J , self-conditionally permissible relative to L . Thus, by Weak Independent Agglomeration, $a_{\bar{J}} \& a_L$ is self-conditionally permissible relative to $J \times L = \{a_J \& a_L, a_J \& a_{\bar{L}}, a_{\bar{J}} \& a_L, a_{\bar{J}} \& a_{\bar{L}}\}$. Hence, by Self-Conditional Alpha, $a_J \& a_{\bar{L}}$ is self-conditionally permissible relative to *Joy or Lesserjoy*, i.e., $\{a_J \& a_{\bar{L}}, a_{\bar{J}} \& a_L\}$, *contra* the improvement claim.

Of course, one could take this tension to amount to a refutation of the procreative asymmetry. But those who are convinced that there really is a procreative asymmetry have reason to be suspicious of the improvement claim. For what it's worth, I'm not convinced that *Joy or Lesserjoy* is a counterexample to strong or stable actualism.

⁴³ Boonin (2014: Sect. 7.3) formalizes the tension by appeal to a transitivity principle. As Boonin points out, however, the transitivity principle he appeals to does not entail that the procreative asymmetry is inconsistent with the improvement claim.

⁴⁴ Weak Independent Agglomeration resembles a principle that Hare (2016: 460) calls "Weak Agglomeration".

⁴⁵ If, as I believe, a_L is not just self-conditionally permissible, but self-conditionally obligatory relative to $L = \{a_L, a_{\bar{L}}\}$, then we can bring out the tension by appeal to an even more plausible principle: namely,

Weaker Independent Agglomeration. If some $a \in A$ is, independently of B , self-conditionally obligatory relative to A , and some $b \in B$ is, independently of A , self-conditionally permissible relative to B , then $a \& b$ is self-conditionally permissible relative to $A \times B$.

⁴⁶ It's not obvious that Weak Independent Agglomeration is compatible with incommensurability; thanks to Caspar Hare for discussion. For criticism of Self-Conditional Alpha, see e.g. Roberts (2003b: 16–40).

10 Completing stable actualism

The last objection to stable actualism targets its incompleteness. In some cases—*Misery or Moremisery* and *Misery or Equalmisery*, for example—every option that maximizes value does so elusively. In those cases, stable actualism goes silent. Unlike strong actualism, stable actualism is incomplete, and one might wonder whether there is any plausible way of completing stable actualism.

Honest answer: I'm not sure. A complete moral theory would need to find some place for the parts of morality that are set aside by the “things being appropriately equal” clause, and would thus need to reckon with some familiar problems in population ethics, including the repugnant conclusion and the mere addition paradox. But there are ways of completing stable actualism that are, I think, clearly superior to strong actualism. So the objection from incompleteness is no obstacle to the modest claims I seek to defend.

The simplest complete theory that entails stable actualism is *hardline actualism*. According to hardline actualism, the antecedent of Stable Maximization is redundant:

Hardline actualism The options that are permissible at world w are all and only the options that stably maximize value at w .

Hardline actualism is simple. It entails that elusive permission is impossible, and thus is consistent with No Weak Dilemmas. It entails stable actualism, and thus can explain both halves of the procreative asymmetry. It handles *Joy or Nothing* and *Misery or Nothing* correctly, and, in *Misery or Moremisery*, it avoids the problematic prediction that we are obligated to create Moremisery if we create Misery.

One objection to hardline actualism is that it conflicts with No Dilemmas. According to hardline actualism, both *Misery or Moremisery* and *Misery or Equalmisery* are cases in which no option is permissible. Another objection to hardline actualism is based in intuition. *Contra* hardline actualism, some people think that, in *Misery or Moremisery*, we are unconditionally obligated to create Misery.

I think that hardline actualism is defensible because I think a strong case can be made that *Misery or Moremisery* is a dilemma. (If we have created Misery, then creating Misery is worse than creating Moremisery for some actual person and not better than creating Moremisery for any actual person.) But my goal is to defend stable actualism, not hardline actualism, so let me turn to another way of completing stable actualism, which is consistent with No Dilemmas and entails that we are, in *Misery or Moremisery*, unconditionally obligated to create Misery.

To get the view on the table, I need introduce one bit of terminology. Recall that $Va_i(a_j)$ is the a_i -value of option a_j . If we let $\operatorname{argmax}_{a_z \in A} Va_i(a_z)$ be the value of an option that maximizes a_i -value, then we can define the *regret* of option a_i as the difference between $Va_i(a_i)$ and $\operatorname{argmax}_{a_z \in A} Va_i(a_z)$. We then have:

Hierarchical actualism. If some option stably maximizes value at w , then the options that are permissible at w are all and only the options that stably maximize value at w . If no option stably maximizes value at w , then the permissible options at w are all and only the options that minimize regret.

The regret of an option is always zero or positive, and the regret of an option that stably maximizes value is zero, so an option that stably maximizes value always minimizes regret. But there are cases in which every option has positive regret: *Misery or Equalmisery* is one; *Misery or Moremisery* is another. According to hardline actualism, an option with positive regret is never permissible. According to hierarchical actualism, an option with positive regret is permissible if the regret of every other option is at least as great. In *Misery or Equalmisery*, the regret of creating Misery is equal to the regret of creating Equalmisery, so hierarchical actualism predicts that both options are unconditionally permissible. In *Misery or Moremisery*, the regret of creating Misery is less than the regret of creating Moremisery, so hierarchical actualism predicts that we are unconditionally obligated to create Misery.

Of course, there are other complete moral theories that entail stable actualism, but these two serve the current argumentative need. As between hardline actualism and strong actualism, I think that hardline actualism is clearly superior, and as between hierarchical actualism and strong actualism, I think that hierarchical actualism is clearly superior.

11 Conclusion

I've argued for two main claims. The first is comparative and modest. I've argued that stable actualism is superior to strong actualism. My reasons are four. Stable actualism can explain the procreative asymmetry,⁴⁷ does not predict that we are obligated to create sure-to-be-even-more-miserable people if we create sure-to-be-miserable people, is consistent with No Weak Dilemmas, and is consistent with impossibility of elusive permission.

The second claim is much less modest. I've argued that there really is an attractive asymmetry: that impermissible options are sometimes self-conditionally permissible, but that permissible options are never self-conditionally impermissible. And I've argued that what gives rise to the attractive asymmetry is an asymmetry in the nature of permission-making.

⁴⁷ Even with regard to the procreative asymmetry, my claims are modest. I have argued that stable actualism can explain the procreative asymmetry; but whether it provides the best explanation remains to be seen. One stout opponent is Roberts' variabilism; see Roberts (2011a). For criticism of Roberts' variabilism, see Algander (2015).

Acknowledgements Thanks to the editors and anonymous reviewers for this journal. I'm particularly grateful to one anonymous reviewer, who helped the paper in a variety of ways. My thanks also, for comments, criticism, and encouragement along the way, to David Builes, Thomas Byrne, Caspar Hare, Michele Odisseas Impagnatiello, Daniel Muñoz, Kieran Setiya, and Brad Skow.

References

- Adler, M. D. (2009). Future generations: A prioritarian view. *The George Washington Law Review*, 77, 1478–1520.
- Ahmed, A. (2014). Dicing with death. *Analysis*, 74, 587–594.
- Algander, P. (2015). Variabilism is not the solution to the asymmetry. *Thought*, 4, 1–9.
- Arrhenius, G. (2009). Can the person affecting restriction solve the problems in population ethics? In M. A. Roberts & D. T. Wasserman (Eds.), *Harming future persons*. Springer.
- Arrhenius, G. (2015). The affirmative answer to the existential question and the person-affecting restriction. In I. Hirose & A. Reisner (Eds.), *Weighing and reasoning: themes from the philosophy of John Broome*. Oxford University Press.
- Arrhenius, G., & Rabinowicz, W. (2010). Better to Be than Not to Be? In H. Joas & B. Klein (Eds.), *The benefits of broad horizons: Intellectual and institutional preconditions for a global social science. Festschrift for Bjorn Wittrock on the occasion of his 65th birthday*. Brill.
- Boonin, D. (2014). *The non-identity problem and the ethics of future people*. Oxford University Press.
- Broome, J. (1999). *Ethics out of economics*. Cambridge University Press.
- Broome, J. (2004). *Weighing lives*. Oxford University Press.
- Broome, J. (2005). Should we value population? *The Journal of Political Philosophy*, 13(4), 399–413.
- Buchak, L. (2017). Taking risks behind the veil of ignorance. *Ethics*, 127, 610–644.
- Bykvist, K. (2006). Prudence for changing selves. *Utilitas*, 18, 364–383.
- Bykvist, K. (2007a). The benefits of coming into existence. *Philosophical Studies*, 135, 335–362.
- Bykvist, K. (2007b). Violations of normative invariance: Some thoughts on shifty oughts. *Theoria*, 73, 98–120.
- Carlson, E. (1995). *Consequentialism reconsidered*. Kluwer.
- Chappell, R. Y. (2017). Rethinking the asymmetry. *Canadian Journal of Philosophy*, 47, 167–177.
- Egan, A. (2007). Some counterexamples to causal decision theory. *Philosophical Review*, 116, 94–114.
- Elstein, D. (2005). The asymmetry of creating and not creating life. *The Journal of Value Inquiry*, 39, 49–59.
- Frick, J. (2020). Conditional reasons and the procreation asymmetry. *Philosophical Perspectives*, 34, 53–87.
- Hare, C. (2007). Voices from another world: Must we respect the interests of people who do not, and will never, exist? *Ethics*, 117, 498–523.
- Hare, C. (2011). Obligation and regret when there is no fact of the matter about what would have happened if you had not done what you did. *Noûs*, 45, 190–206.
- Hare, C. (2016). Should we wish well to all? *Philosophical Review*, 125, 451–472.
- Hare, R. M. (1975). Abortion and the golden rule. *Philosophy and Public Affairs*, 4, 201–222.
- Harman, E. (1999). Creation ethics: The moral status of early fetuses and the ethics of abortion. *Philosophy and Public Affairs*, 28, 310–324.
- Harman, E. (2004). Can we harm and benefit in creating? *Philosophical Perspectives*, 18, 89–113.
- Harman, E. (2009). “I’ll Be Glad I Did It” reasoning and the significance of future desires. *Philosophical Perspectives*, 23, 177–199.
- Harper, W. (1986). Mixed strategies and ratifiability in causal decision theory. *Erkenntnis*, 24, 25–36.
- Heyd, D. (1992). *Genethics: Moral issues in the creation of people*. University of California Press.
- Holtug, N. (2001). On the value of coming into existence. *The Journal of Ethics*, 5, 361–384.
- Howard-Snyder, F. (2008). Damned if you do; damned if you don’t. *Philosophia*, 36, 1–15.
- Johansson, J. (2010). Being and betterness. *Utilitas*, 22, 285–302.
- Kavka, G. (1982). The paradox of future individuals. *Philosophy and Public Affairs*, 11, 93–112.
- Kotzen, M. (MS). Lewis’s argument against desire-as-belief.
- McConnell, T. (2018). Moral dilemmas. *The Stanford Encyclopedia of Philosophy* (Fall 2018), E. Zalta (ed.). Retrieved <<https://plato.stanford.edu/archives/fall2018/entries/moral-dilemmas/>>.
- McMahan, J. (1981). Problems of population theory. *Ethics*, 92, 96–127.

- McMahan, J. (1994). Review of genetics: Moral issues in the creation of people. *Philosophical Review*, 103, 557–559.
- McMahan, J. (2009). Asymmetries in the morality of causing people to exist. In M. A. Roberts & D. T. Wasserman (Eds.), *Harming future people*. Springer.
- Narveson, J. (1967). Utilitarianism and new generations. *Mind*, 76, 62–72.
- Narveson, J. (1973). Moral problems of population. *The Monist*, 57, 62–86.
- Narveson, J. (1978). Future people and us. In R. I. Sikora & B. Barry (Eds.), *Obligations to future generations*. White Horse Press.
- Parfit, D. (1982). Future generations: Further problems. *Philosophy and Public Affairs*, 11, 114–172.
- Parfit, D. (1987). *Reasons and persons*. Oxford University Press.
- Persson, I. (2009). Rights and the asymmetry between creating good and bad lives. In M. A. Roberts & D. T. Wasserman (Eds.), *Harming future persons*. Springer.
- Prichard, H. A. (1968). *Moral obligation, duty and interest: Essays and lectures by H.A. Prichard*. Oxford University Press.
- Roberts, M. A. (2003a). Can it ever have been better never to have existed at all? Person-based consequentialism and a new repugnant conclusion. *Journal of Applied Philosophy*, 20, 159–185.
- Roberts, M. A. (2003b). Is the person-affecting intuition paradoxical? *Theory and Decision*, 55, 1–44.
- Roberts, M. A. (2010). *Abortion and the moral significance of merely possible persons: Finding middle ground in hard cases*. Springer.
- Roberts, M. A. (2011a). The asymmetry: A solution. *Theoria*, 77, 333–367.
- Roberts, M. A. (2011b). An asymmetry in the ethics of procreation. *Philosophy Compass*, 6, 765–776.
- Roberts, M. A. (2019). The nonidentity problem. *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), E. Zalta (ed.). Retrieved from <<https://plato.stanford.edu/archives/sum2019/entries/nonidentity-problem/>>.
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffman (Eds.), *Modality: Metaphysics, logic, and epistemology*. Oxford University Press.
- Sen, A. (1970). *Collective choice and social welfare*. Holden-Day.
- Setiya, K. (2014a). What is a reason to act? *Philosophical Studies*, 167, 221–235.
- Setiya, K. (2014b). The ethics of existence. *Philosophical Perspectives*, 28, 291–301.
- Shiffrin, S. (1999). Wrongful life, procreative responsibility, and the significance of harm. *Legal Theory*, 5, 117–148.
- Spencer, J., & Wells, I. (2019). Why take both boxes? *Philosophy and Phenomenological Research*, 99, 27–48.
- Spencer, J. (forthcoming). Rational monism and rational pluralism. *Philosophical Studies*.
- Sterba, J. P. (1987). Explaining asymmetry: A problem for parfit. *Philosophy and Public Affairs*, 16, 188–192.
- Wedgwood, R. (2015). The pitfall of ‘reasons.’ *Philosophical Issues*, 25, 123–143.
- Williams, B. (1981). Internal and external reasons. In B. Williams (Ed.), *Moral luck*. Cambridge University Press.
- Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford University Press.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.