

The Procreative Asymmetry and the Impossibility of Elusive Permission

Jack Spencer

Words: 7,801

Words, including footnotes: 9,022

This essay is about two asymmetries: the procreative asymmetry and what I will call the “attractive asymmetry.” I seek to defend and explain the attractive asymmetry, and to use the attractive asymmetry to develop a way of explaining the procreative asymmetry.

1/ Introduction

The procreative asymmetry is a puzzle in the ethics of creation that stems from a pair of hard-to-reconcile intuitions.¹

According to the first intuition, not creating sure-to-be-happy people is permissible. 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.

If we do not create Joy, it seems that we’ve done something morality permits. And this intuition does not depend on the usual trappings of creation. We may suppose that we can create Joy by pressing a button, and that Joy, if created, will never causally interact with us. Even with these additional suppositions, it seems clear that we have done something morality permits if we do not create Joy.

According to the second intuition, not creating sure-to-be-miserable people is obligatory. 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.

If we do not create Misery, it seems that we’ve done something morality requires.

¹ 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 (MS), 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).

Taken together, however, these two intuitions create an explanatory puzzle. Whence the asymmetry? If the misery that would have filled Misery's life makes not creating Misery obligatory, why doesn't the joy that would have filled Joy's life make not creating Joy impermissible? We could deny one or both of the intuitions, of course. But if we accept both intuitions, accepting that there really is a procreative asymmetry—as I will, hereafter—then we need an explanation.²

My proposed explanation amends an existing proposal. In an attempt to explain the procreative asymmetry, many consequentialists have been drawn to *actualism*,³ a moral theory that takes inspiration from Narveson's dictum: *make people happy, not happy people*.⁴ Unfortunately, actualism delivers only half of what it promises. It's capable of explaining the happy half of the procreative asymmetry, but it's not capable of explaining the miserable half.⁵ And as its critics have pointed out, it admits of other counterexamples, as well.⁶

But the failure of actualism should not convince us to abandon its motivation. Actualism is motivated by two intuitions. The first is the *person-affecting intuition*: that nothing can be better or worse unless it's better or worse for someone.⁷ The second is the *comparability intuition*: that happiness is better for someone than nonexistence, which in turn is better for them than misery. Actualism is one theory that captures these intuitions, but it's not the most plausible theory that does so. Those who find these intuitions compelling should reject actualism and instead accept some moral theory that entails what I call "stable actualism."

Stable actualism is principled. It avoids the clearest counterexamples to actualism. And it has another virtue, too: it can explain both halves of the procreative asymmetry. Stable actualism is not for everyone. Those who do not find the motivation for actualism compelling are unlikely to accept it. But even the staunchest opponents of actualism can

² The procreative asymmetry is sometimes stated in terms of reasons; see *e.g.* Chappell (2017: 167), Frick (MS), and McMahan (1981: 100; 2009: 49). Like Wedgwood (2015), I'm somewhat skeptical of "reasons"-talk, so I prefer to state it in deontic terms.

³ See *e.g.* Narveson (1967; 1973; 1978) and Parsons (2002).

⁴ Narveson (1973: 80).

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

⁶ For criticisms of actualism, see *e.g.* Arrhenius (2009; 2015), Bykvist (2007b), Frick (MS), Hare (2007), McMahan (1981; 1994), Roberts (2010: ch. 2; 2011a; 2011b), and Parfit (1982; 1987).

⁷ The person-affecting intuition can—and probably should—be weakened to allow for considerations of equality and priority. Perhaps it could be stated as follows: nothing can be better or worse unless (a) it is better or worse for someone or (b) it has some non-negligible effect on global values, like equality or priority.

accept the modest thesis that I hope to defend: that stable actualism is superior to actualism.

The route from actualism to stable actualism is paved by the other asymmetry mentioned above—the attractive asymmetry. In contrast to the procreative asymmetry, which concerns two first-order ethical intuitions, the attractive asymmetry is a bit of high theory, which resides, not at the ethical level, but at the metaethical level. To get the attractive asymmetry on the table, let me introduce some terminology.

Say that a normative theory gives rise to *normative variance* if it predicts that whether an option is permissible sometimes depends on whether the option is chosen.⁸ A normative theory that gives rise to normative variance can give rise to two sorts of modified permission. On the one hand there can be *attractive permission*: an option is permissible, but wouldn't be if it were unchosen. On the other hand there can be *elusive permission*: an option is permissible, but wouldn't be if it were chosen.

The attractive asymmetry says that attractive permission is possible and that elusive permission is not. Of course, that's controversial. Some philosophers believe that both attractive and elusive permission are possible,⁹ and some believe that both are impossible.¹⁰ But I accept the attractive asymmetry and, in what follows, will defend and explain it.

I will then use the attractive asymmetry to amend actualism. Actualism gives rise to both attractive and elusive permission.¹¹ The fact that it gives rise to attractive permission is a virtue, I think, since it's part of what makes actualism capable of explaining the happy half of the procreative asymmetry. But the fact that it gives rise to elusive permission is, I think, its fundamental defect.

The defect can be remedied, however. We can rid actualism of elusive permission, without ridding it of attractive permission, and the result is stable actualism.

2/ Actualism

⁸ See Carlson (1995: ch 6), who credits both the idea and the name 'normative variance' to Wlodek Rabinowicz. Also see Prichard (1986: 37) and Bykvist (2007b: 99).

⁹ Bykvist (2007a) lists many normative theories that are committed to the possibility of both attractive and elusive permission. As we'll see in section 8, the claim that an agent is permitted to believe any proposition that is likely on her evidence also entails that both attractive and elusive permission are possible.

¹⁰ See *e.g.* Broome (2004: 74), Carlson (1995: ch. 6), Frick (MS), Hare (2007; 2011: n. 11), and Roberts (2010: 62).

¹¹ Hare (2007) distinguishes a strong and weak form of actualism. I'm concerned with the strong form.

Actualism is a perspectival form of consequentialism, which allows the value of an option to depend on which option is chosen. To get a sense of actualism, it's helpful to introduce some machinery.

Let $A = \{a_1, \dots, a_n\}$ be set of the options available to the agent.¹² 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 thus 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 is chosen—then we can visualize the perspectival values by constructing the following $n \times n$ matrix:

	a_1	...	a_n
a_1	$Va_1(a_1)$...	$Va_1(a_n)$
...
a_n	$Va_n(a_1)$...	$Va_n(a_n)$

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.

It will be helpful to make the intuitions that motivate actualism a bit more precise.

To capture the person-affecting intuition, I propose a Pareto principle.¹³ For each $a \in A$, let P_a be the population that would exist if a were chosen. We then can formulate the following three-part principle:

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

- (i) If a_i is better than a_j for someone in P_a and not worse for anyone in P_a , then, $Va(a_i) > Va(a_j)$.
- (ii) If a_i is worse than a_j for someone in P_a and not better for anyone in P_a , then $Va(a_i) < Va(a_j)$.
- (iii) If a_i is neither better nor worse than a_j for anyone in P_a , then $Va(a_i) = Va(a_j)$.

¹² 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 always mutually exclusive and jointly exhaustive.

¹³ Here I follow Hare (2007: 502).

Perspectival Pareto is oversimplified. It's not sensitive to global values, like equality or priority.¹⁴ But for the questions in the ethics of creation that I want to explore here, I think that we can safely set those sorts of global values aside. So instead of operating with some more complicated principle, I will operate with Perspectival Pareto.

To capture the comparability intuition, I propose the following principle:

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.¹⁵

Comparability is controversial because it conflicts with two prima facie plausible claims. Let x be some actual person who leads a miserable life. 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 $w_{@}$.¹⁶

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 $w_{@}$.

But this argument against Comparability is not convincing, because, as Holtug (2001) and Arrhenius and Rabinowicz (2010) argue, Accessibility is false. An analogy with preference might be helpful. An actual person can prefer a world in which they never exist

¹⁴ The conflict between 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.

¹⁵ If people are necessary beings, as necessitists, like Williamson (2013), maintain, then we should replace talk of existence with talk of chunkiness.

¹⁶ 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).

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 for goodness, I think: $w_{@}$ 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 .

The third assumption of actualism is a principle that connects values to deontic status:

Maximization. An option is permissible at world w if and only if it maximizes value at w .

Options that maximize value at the actual world will be said to maximize value, *sans phrase*.

Maximization might seem to go without saying, but I think it's the crucial principle. As we'll see, stable actualists accept Perspectival Pareto and Comparability, but reject Maximization.

Before getting into stable actualism, however, let me say a bit more about actualism.

3/ Actualism and Joy

One advantage of actualism is that it can explain the happy half of the procreative asymmetry.

According to actualism, *Joy or Nothing* exhibits normative variance: if we create Joy, then we are obligated to create Joy; 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*, according to actualism, because creating Joy is better than not creating Joy for someone who would exist if we created Joy and not worse for anyone who would exist if we created Joy. Not creating Joy is *self-conditionally permissible*, according to actualism, because not creating Joy is neither better nor worse for anyone who would exist if we did not create Joy.

The happy half of the procreative asymmetry is thus explained. According to actualism, if we do not create Joy, then what we do is permissible because not creating Joy is neither better nor worse for anyone.

Actualism also provides a neat explanation of the 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. If we create Joy, however, the only option that retrospectively appears choiceworthy is creating Joy; the fitting retrospective attitude is gladness. There are many ways one could try to explain this retrospective asymmetry, but the explanation actualism offers is particularly appealing. According to actualism, the retrospective asymmetry is just a reflection of the normative

variance. The options that retrospectively appear choiceworthy are exactly the permissible options.

Actualism faces challenges, however. There are four main objections, by my count. One objection applies to any theory that gives rise to normative variance, so I will discuss it later. The other three objections are specific to actualism, so I will consider them forthwith.

4/ Actualism and Lesserjoy

The first objection concerns nonidentity cases.¹⁷ Some critics think that actualism delivers the wrong results in cases like the following:

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 a 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 actualism, *Joy or Lesserjoy* exhibits normative variance: we are obligated to create whoever we create. But some philosophers defend *the improvement claim*: that we are *unconditionally* obligated to create Joy, i.e., obligated to create Joy, no matter which option we choose.¹⁸ And if the improvement claim is true, actualism is false.

The improvement claim does not threaten the modest thesis I seek to defend. Actualism and stable actualism are both 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. 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 to choosing.¹⁹ 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.

¹⁷ The literature on the nonidentity problem, and the conflict between it and person-affecting intuitions, is vast. See Roberts (2019) and the references therein.

¹⁸ 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 of equality and priority, which I have set aside.

¹⁹ For arguments that retrospective gladness and permissibility do not co-vary, see Harman (2009) and Setiya (2014b).

If we create Lesserjoy, we should not wish that we had created Joy; the fitting retrospective attitude is gladness. So if the improvement claim is true, there will be cases in which we should be glad to have chosen an option that we knew to be impermissible prior to choosing. (Actualists, by contrast, can *accept* that retrospective gladness and permissibility co-vary, since the retrospective asymmetry in *Joy or Lesserjoy* exactly parallels the normative variance that actualism alleges to obtain.)

Another concern comes from the tension between the improvement claim and the procreative asymmetry.²⁰

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. (An option is self-conditionally permissible if it is permissible if chosen. One who rejects normative variance will think that every self-conditionally permissible option is unconditionally permissible.) But the claim that both options in both choices are self-conditionally permissible is hard to square with the improvement claim. If it is permissible to not create Joy (in *Joy or Nothing*) and 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*).

One way to formalize this tension appeals to two principles.²¹ Let A and B be two sets of options. As a matter of terminology, 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 principle:²²

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-

²⁰ For other discussions of the interaction between the improvement claim and the procreative asymmetry, see *e.g.* Boonin (2014: ch. 7), Frick (MS), Harman (2004), Heyd (1992), McMahan (1981), Narveson (1973; 1978), Roberts (2010; 2011a; 2019) and Parfit (1987).

²¹ 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.

²² Weak Independent Agglomeration resembles a principle that Hare (2016: 460) calls "Weak Agglomeration."

conditionally permissible relative to B , then $a \& b$ is self-conditionally permissible relative to $A \times B$.²³

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,²⁴ but both are plausible. And given the two principles, the improvement claim is inconsistent with the procreative asymmetry. Let $J = \{Joy, \sim Joy\}$ be the set that contains the options of creating and not creating Joy, and let $L = \{Less, \sim Less\}$ 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, $\sim Joy \& Less$ is self-conditionally permissible relative to $J \times L = \{Joy \& Less, Joy \& \sim Less, \sim Joy \& Less, \sim Joy \& \sim Less\}$. Hence, by Self-Conditional Alpha, $\sim Joy \& Less$ is self-conditionally permissible relative to $\{Joy \& \sim Less, \sim Joy \& Less\}$, *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 am not convinced that *Joy or Lesserjoy* is a counterexample to actualism.

5/ Actualism and Misery

²³ If, as I believe, *Less* is not just self-conditionally permissible, but self-conditionally obligatory relative to $L = \{Less, \sim Less\}$, 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 $A \times B$.

²⁴ It's not obvious that Weak Independent Agglomeration is compatible with incommensurability; thanks to [redacted] for discussion. For criticism of Self-Conditional Alpha, see *e.g.* Roberts (2003b: 16-40).

I am convinced that there are counterexamples to actualism, however. In fact, we've seen one already: *Misery or Nothing*. According to actualism, *Misery or Nothing* exhibits normative variance: if we create Misery, then we are obligated to not create Misery; if we do not create Misery, then both options are permissible. In making these predictions, actualism errs twice. It wrongly predicts that *Misery or Nothing* exhibits normative variance, and it wrongly predicts that we are permitted to create Misery if we don't. Actualism often accords with our retrospective attitudes, but, strikingly, in *Misery or Nothing*, it doesn't. Unlike what actualism predicts, creating Misery retrospectively appears regrettable, no matter which option we choose.

And things get worse. For once we've seen one counterexample, we can 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 am open to the view that it's a dilemma, in which both options are unconditionally impermissible, and I am open to the view that creating Misery is unconditionally obligatory. But I am not open to what actualism says. According to actualism, *Misery or Moremisery* exhibits normative variance: we are obligated to create whoever we do not create. And this prediction of actualism—that we are obligated to create Moremisery if we create Misery—seems clearly false.²⁵

6/ Actualism and Deontic Consistency

²⁵ Roberts (2010: ch.2) argues that two other examples make trouble for 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 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.

A third objection to actualism also succeeds. Actualism violates deontic consistency principles that should be affirmed.

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

No Dilemmas. It's impossible that an agent face a dilemma.

No Dilemmas is controversial. It's opponents often point to Sophie's choice scenarios, in which an agent must decide which of two innocent people will suffer some horrible fate.²⁶ But an even stronger challenge to No Dilemmas comes from symmetric miserable creation cases. Consider:

Misery or Equalmisery. We are 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 and Equalmisery will be affected by our choice.

In Sophie's choice scenarios, the option that is chosen is better for *someone*: namely, the one who does not suffer the horrible fate. In *Misery or Equalmisery*, the option that is chosen is better for *no-one*, and much, much worse for whoever we create. I think that a strong case can be made that *Misery or Equalmisery* is a genuine dilemma.

No Dilemmas is not, at this point in the dialectic, however, the relevant deontic consistency principle. Actualism is consistent with No Dilemmas.

An agent faces 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 actualism is that it conflicts with:

No Weak Dilemmas. It's impossible that an agent face a weak dilemma.

The conflict is brought out by *Misery or Equalmisery*. According to 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 odd, to put it mildly. Here's Bykvist bringing out the oddity:

²⁶ See McConnell (2018) and references therein.

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 unavailable 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 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 actualism and No Weak Dilemmas as a strong reason to reject actualism, and I think they're right. I think No Weak Dilemmas should be affirmed.

7/ Stable Actualism

Since actualism admits of counterexamples, at least one of its three core assumptions—Perspectival Pareto, Comparability, or Maximization—is false. Comparability is true, and anyway not to blame for the counterexamples. So the choice is between Perspectival Pareto and Maximization.

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. 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.²⁷ But the alleged distinction between “positive” and “negative” value is mysterious, as it 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.

There is a pattern to the counterexamples above. In all of them, an option that actualism predicts to be self-conditionally impermissible is wrongly predicted by actualism to be permissible if unchosen. This suggests the principle that should replace Maximization.

We can distinguish two ways that an option might maximize value. If *a* maximizes value and also maximizes *a*-value, then we’ll say that *a* *stably* maximizes value. If *a* maximizes value but does not maximize *a*-value, then we’ll say that *a* *elusively* maximizes value. Actualism is blind to the distinction between stable and elusive maximization. But the cases above suggest that the distinction is important, and that we should replace Maximization with:

Stable Maximization. If some option stably maximizes value at *w*, then an option is permissible at *w* if and only if it stably maximizes value at *w*.

Let *stable actualism* be the view that results from combining Perspectival Pareto, Comparability, and Stable Maximization.

If every option that maximizes value also stably maximizes value, then actualism and stable actualism coincide. Actualism and stable actualism thus coincide in *Joy or Nothing* and in *Joy or Lesserjoy*.

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 actualism, stable actualism is incomplete. Later I will explore two ways of completing stable actualism. But there is something interesting and important to point out about stable actualism, itself, incomplete though it is. *Stable actualism can explain both halves of the procreative asymmetry*.

Consider *Misery or Nothing*, and let *Misery* and \sim *Misery* be the options of creating and not creating Misery, respectively. As we saw above:

- (1) Both *Misery* and \sim *Misery* maximize \sim *Misery*-value.
- (2) Only \sim *Misery* maximizes *Misery*-value.

²⁷ This tempting view is akin to the one Parfit (1987: 525-6) regards as providing the “best” explanation of the procreative asymmetry.

Two other claims are also true:

(3) If we do not create Misery, then: \sim Misery stably maximizes value just if it maximizes \sim Misery-value, and Misery stably maximizes value just if it maximizes both \sim Misery-value and Misery-value.

(4) If we create Misery, then: \sim Misery stably maximizes value just if it maximizes both Misery-value and \sim Misery-value, and Misery stably maximizes value just if it maximizes Misery-value.

Claims (1)-(4) entail that, no matter what we do, the only option that stably maximizes value is \sim Misery. So, unlike actualism, stable actualism can explain *both* halves of the procreative asymmetry. It correctly handles *Joy or Nothing*, and it also correctly handles *Misery or Nothing*, predicting that we are unconditionally obligated not to create Misery.

Over the next three sections, I will consider three objections to stable actualism. The first alleges that stable actualism is *ad hoc*. The second—the objection that was previously postponed—alleges that stable actualism, like any theory that gives rise to normative variance, should be rejected. The third alleges that there is no plausible way to complete stable actualism.

8/ 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 .²⁸

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.

²⁸ This principle resembles a principle that Hare (2011: 196) calls, "Reasons are not Self-Undermining." It also resembles a principle that [redacted].

With Stability in view, we can better understand what's going on in *Misery or Nothing*. If we do not create Misery, then creating Misery maximizes value. But creating Misery would not maximize value if we created Misery. So, if Stability is true, creating Misery cannot be made permissible by the fact that it maximizes value, *contra* Maximization.

I believe that Stability holds true, 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 the agent's evidence, then 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. Let p be the proposition that I would be a good politician.²⁹ 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 very 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 evidential probability that p enjoys.

²⁹ Here I adapt an example that Kotzen (MS) uses in his discussion of desire-as-belief.

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 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. I think 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.³⁰

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.³¹

The second ancillary 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, then, is 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

³⁰ 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), and [redacted].

³¹ See *e.g.* Rosen (2010: 118).

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 hold, it follows that elusive permission is impossible.

I think that all permissions are derivative, so I think Derivative Elusive Permissions follows from a more general truth. 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.

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*.

9/ The Possibility of Normative Variance

Like actualism, stable actualism entails that normative variance is possible. Some philosophers think that normative variance is not possible.³²

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)

If normative variance is impossible, then stable actualism is false.

But why think that normative variance is impossible? What's implausible about it?³³

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

³² Also see *e.g.* Hare (1975: 219) and Narveson (1978: 44).

³³ Both Bykvist (2007b) and Howard-Snyder (2008) defend the possibility of normative variance.

theories that give rise to normative variance conflict with these principles. I'm skeptical of No Dilemmas, as I said. But I think that 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 that 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 the next section, I'll offer one way of completing stable actualism that is 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 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)³⁴

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, just by deliberating, winnow her options to just the permissible ones. But the agent *can*, just by deliberating, winnow her options 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

³⁴ Here Bykvist echoes Carlson (1995: 101).

agent *can*, just by deliberating, winnow her options 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, just 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 faces Buridan's ass, 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 have not already decided whether we will create Joy, it takes an act of will, over and above sound deliberation, to choose an option.

I am not aware of other arguments against the possibility of normative variance, but a number of people who have seen this paper have wondered at this point about the relationship between normative variance and practical reasoning.³⁵ 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 .³⁶ If permission-makers are reasons and this connection between reasons and reasoning holds,³⁷ 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 .

And 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.³⁸ A sound bit of practical reasoning has no false premises. So, if p does not hold at the world 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

³⁵ Thanks to [redacted] for discussion.

³⁶ See *e.g.* Hare (2011), Setiya (2014a), and Williams (1981).

³⁷ It's not obvious that permission-makers are reasons; thanks to [redacted] for discussion.

³⁸ *Cf.* Hare (2011: 196; 2011: n. 11).

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. We have already seen three potential illustrations of attractive permission: *Joy or Nothing*, *Joy or Lesserjoy*, 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 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.

10/ Completing Stable Actualism

The last objection to stable actualism targets its incompleteness. *Question:* Is there any plausible complete moral theory that entails stable actualism?

Answer: I'm not sure. A complete moral theory would need to find some place for the global values that I have set aside, and that would require coming to terms 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 actualism. So the objection from incompleteness is no obstacle to the modest thesis I seek to defend.

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

Hardline Actualism. An option is permissible at world w if and only if it stably maximizes value at w .

Hardline actualism is simple. It entails that elusive permission is impossible and thus is consistent with No Weak Dilemmas. It vindicates both the person-affecting intuition and the comparability intuition. It entails stable actualism, and thus can explain both halves of the procreative asymmetry. 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, *Misery or Moremisery* and *Misery or Equalmisery* are both dilemmas, in which neither 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 that a strong case can be made that *Misery or Moremisery* is a dilemma. (After all, if we create Misery, then creating Misery is better for no-one and much, much worse for Misery.) But my goal is to defend stable actualism, not hardline actualism, so let me turn to another way of completing actualism, which is consistent with No Dilemmas and entails that we are unconditionally obligated to create Misery, in *Misery or Moremisery*.

To get the view on the table, I need introduce one bit of terminology. Recall that $V_{a_j}(a_i)$ is the a_j -value of option a_i , i.e., the value option a_i , if option a_j is chosen. If we let $\operatorname{argmax}_{a_z \in A} V_{a_i}(a_z)$ 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 $V_{a_i}(a_i)$ and $\operatorname{argmax}_{a_z \in A} V_{a_i}(a_z)$. We can then state:

Hierarchical Actualism. If some option stably maximizes value at w , then the permissible options at w are the options that stably maximize value at w . If no option stably maximizes value at w , then the options that are unconditionally permissible at w are 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 actualism, I think that hardline actualism is clearly superior, and as between hierarchical actualism and 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 actualism. My reasons are four: (1) unlike actualism, stable actualism can explain the procreative asymmetry;³⁹ (2) unlike actualism, stable actualism does not predict that we are obligated to create sure-to-be-even-more-miserable people if we create sure-to-be-miserable people; (3) unlike actualism, stable actualism is consistent with No Weak Dilemmas; and (4) unlike actualism, stable actualism is consistent with the claim that elusive permission is impossible.

The second claim is much less modest. I've argued that there really is an attractive asymmetry. Attractive permission is possible: an impermissible option can be self-conditionally permissible. But elusive permission is not possible: a permissible option cannot be self-conditionally impermissible. And I've argued that what gives rise to the attractive asymmetry is an asymmetry in the nature of permission-making.⁴⁰

References

Adler, M. D. (2009). Future Generations: A Prioritarian View. *The George Washington Law Review* 77: 1478-520.

³⁹ 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).

⁴⁰ [Acknowledgements]

- Ahmed, A. (2014). Dicing with Death. *Analysis* 74: 587-94.
- 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 and D. T. Wasserman (eds.), *Harming Future Persons*, pp. 289-314. Springer.
- (2015) The Affirmative Answer to the Existential Question and the Person-Affecting Restriction. In I. Hirose and A. Reisner (eds), *Weighing and Reasoning: Themes from the Philosophy of John Broome*, pp. 110-28. Oxford University Press.
- Arrhenius, G and Rabinowicz, W. (2010). Better to Be than Not to Be? In H. Joas and 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*, pp. 399-414. 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.
- (2004). *Weighing Lives*. Oxford University Press.
- (2005) Should We Value Population? *The Journal of Political Philosophy* 13.4: 399-413.
- Bykvist, K. (2007a). The Benefits of Coming into Existence. *Philosophical Studies* 135: 335-62.
- (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: 167-77.
- 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. (MS). Conditional Reasons and the Procreation Asymmetry.
- 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.
- (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.
- (2016). Should We Wish Well to All? *Philosophical Review*: 125: 451-72.
- Hare, R. M. (1975). Abortion and the Golden Rule. *Philosophy and Public Affairs* 4: 201-22.
- Harman, E. (1999). Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion. *Philosophy and Public Affairs* 28: 310-24.

- (2004). Can We Harm and Benefit in Creating? *Philosophical Perspectives* 18: 89-113.
- (2009). "I'll Be Glad I Did It" Reasoning and the Significance of Future Desires. *Philosophical Perspectives* 23: 177-99.
- 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: 361-84.
- Howard-Snyder, F. (2008). Damned if You Do; Damned If You Don't. *Philosophia* 36: 1-15.
- 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.), URL = <<https://plato.stanford.edu/archives/fall2018/entries/moral-dilemmas/>>.
- McMahan, J. (1981). Problems of Population Theory. *Ethics* 92: 96-127.
- (1994). Review of *Genethics: Moral Issues in the Creation of People*. *Philosophical Review* 103: 557-9.
- (2009). Asymmetries in the Morality of Causing People to Exist. In M. A. Roberts and D. T. Wasserman (eds.) *Harming Future People*, pp. 49-68. Springer.
- Narveson, J. (1967). Utilitarianism and New Generations. *Mind* 76: 62-72.
- (1973). Moral Problems of Population. *The Monist* 57: 62-86.
- (1978). Future People and Us. In *Obligations to Future Generations*, R. I. Sikora and B. Barry (eds.), pp. 38-60. White Horse Press.
- Parfit, D. (1982). Future Generations: Further Problems. *Philosophy and Public Affairs* 11: 114-72.
- (1987). *Reasons and Persons*. Oxford University Press.
- Parsons, J. (2002). Axiological Actualism. *Australasian Journal of Philosophy* 80: 137-47.
- Persson, I. (2009). Rights and the Asymmetry Between Creating Good and Bad Lives. In M. A. Roberts and D. T. Wasserman (eds.) *Harming Future Persons*, pp. 29-47. 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: 159-85.
- (2003b). Is the Person-Affecting Intuition Paradoxical? *Theory and Decision* 55: 1-44.
- (2010). *Abortion and the Moral Significance of Merely Possible Persons: Finding Middle Ground in Hard Cases*. Springer.

- (2011a). The Asymmetry: A Solution. *Theoria* 77: 333-367.
- (2011b). An Asymmetry in the Ethics of Procreation. *Philosophy Compass* 6: 765-76.
- (2019). The Nonidentity Problem. *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), E. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/sum2019/entries/nonidentity-problem/>](https://plato.stanford.edu/archives/sum2019/entries/nonidentity-problem/).
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In B. Hale and A. Hoffman (eds.), *Modality: Metaphysics, Logic, and Epistemology*, pp. 109-36. Oxford University Press.
- Shiffrin, S. (1999). Wrongful Life, Procreative Responsibility, and the Significance of Harm. *Legal Theory* 5: 117-48.
- Sen, A. (1970). *Collective Choice and Social Welfare*. Holden-Day.
- Setiya, K. (2014a). What Is a Reason to Act? *Philosophical Studies* 167: 221-35.
- (2014b). The Ethics of Existence. *Philosophical Perspectives* 28: 291-301.
- Sterba, J. P. (1987). Explaining Asymmetry: A Problem for Parfit. *Philosophy and Public Affairs* 16: 188-92.
- Wedgwood, R. (2015). The Pitfall of 'Reasons'. *Philosophical Issues* 25: 123-43.
- Williams, B. (1981). Internal and External Reasons. In his *Moral Luck*, pp. 101-13. Cambridge University Press.
- Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford University Press.
- [redacted]