



ARTICLES

The Sequence Argument Against the Procreation Asymmetry

Matthew Adelstein 

University of Michigan, Ann Arbor, MI, USA
Email: Untrappedzoid@gmail.com

Abstract

The procreation asymmetry is a widely held view in ethics, claiming that one should make existing people happy but has no reason to make happy people. Here, I shall present a new objection demonstrating from modest premises that one has a reason to take a sequence of actions that simply creates a happy person; yet this judgment in combination with plausible principles about sequences of actions entails that one has some reason to simply create a happy person. Additionally, I will argue that one's reasons to create a happy person are quite strong.

Keywords: Procreation asymmetry; population ethics; person-affecting view

1. Introduction

The procreation asymmetry, according to which one should make existing people happy but does not have a reason to create happy people is arguably an intuitive one.¹ Few of us feel we have any obligation to have children, even if we expect those children to live good lives.² Various people have devised person-affecting views of population ethics to make sense of this intuition, according to which one has no reason to create happy people (e.g., Frick 2020; Meacham 2012; Wolf 1996). However, the asymmetry has proved surprisingly difficult to vindicate or explain and has come under fire from many.³ In this paper, I will argue that there is no procreation asymmetry; it is good to make happy people, not just to make existing people happy.

Narveson's (1973, p.80) original formulation of the asymmetry as "we are in favor of making people happy, but neutral about making happy people," is slightly ambiguous but seems to suggest that one has no reason to create happy people. Frick (2020, p. 56) summarizes it, saying, "there is no moral reason to create a person whose life would foreseeably

¹In this paper, I use happy in a very expansive sense, to mean whatever makes a person's life go well. I use it instead of the more traditional well-being because previous formulations of the asymmetry have tended to describe the creation of happy people.

²It's worth noting what Chappell (2017) points out, that while most of us think intuitively that there is no obligation to create happy people, it does seem intuitively plausible that we have at least some reason to do so. Thus, it's not clear that such intuitions genuinely lend support to the procreation asymmetry.

³Examples include Parfit (1984), Chappell (2017), and Beckstead (2013).

be worth living, just because her life would be worth living.” Grill (2017, p. 219) claims “we seem to have strong moral reason to improve the wellbeing of existing lives, but, again, no moral reason to create lives that are worth living.” My argument, if successful, would disprove the asymmetry, given that if one should create a happy person at no cost, then there is clearly some reason to create happy people. Note here that I am not concerned with showing that creating happy people is obligatory, merely that one has some reason to do it.

Once we have gotten clear about this definition of the asymmetry, according to which one has no reason at all to create happy people, it is not even clear that the asymmetry is intuitive.⁴ It is not hard to believe that the fact that my parents’ action to give birth to me allowed me to live a flourishing and happy life gave them some moral reason to do it. As Pummer⁵ (2016, p. 237) notes:

... suppose I could push a button that would create billions of happy people living on several large and lush Eden-like planets. These people would in turn produce further generations of happy people, who would do likewise, and so on for the foreseeable future. Pushing the button would cost me nothing and do no harm or wrong. Would it be wrong of me not to push the button, in this case? Yes, I believe it would.

2. The sequence argument

2.1 The argument explained

My argument against the asymmetry appeals to the following three principles each of which are intuitively plausible, and will be defended more later:

Generative Improvement: If one can press a button that will increase the utility of an existing person by some amount and create some person with positive utility, then they should do so, if they will later have the option described in Modification Improvement.

Modification Improvement: If one can, after pressing the button described in the Generative Improvement principle, press another button, which will rescind the utility increase for the existing person, while increasing the utility of the guaranteed future person created by the first button by an amount that is many times larger, they should do so.

Sequential Desirability: If one should press a sequence of buttons in some order, then if another button has the same effects as pressing the buttons in that sequence in that order, then that other button should be pressed.

Together, these principles entail that it is good, *ceteris paribus*, to create a person with some positive utility. To see this, Generative Improvement entails that one should press some button that creates a person with some amount of utility and increases an existing person’s utility by some amount. Then, Modification Improvement entails that one should press a button that would increase the utility of the new person by some

⁴For a detailed discussion, see Chappell (2017).

⁵Notably, Pummer no longer endorses this argument and now defends the procreation asymmetry, but his argument has force nonetheless.

great amount while rescinding the comparatively slight benefit to the existing person. Together then, these two buttons create a person with some positive amount of utility. Thus, a third button which does the same things as pressing the two aforementioned buttons, is worth pressing, all else equal. But this third button just creates a person with positive utility. Thus, all else equal, one should create a person with some amount of positive utility, and the asymmetry is false.

We can see an example of the basic argument in the following scenario. Suppose that one has two buttons. The first button would create a happy person who would live 70 years and would also give a random 70-year-old stranger an extra year of life. The second button would cause the first person to live 90 years instead of 70 but would rescind the extra year of life that would have been given by the first button. Each button is worth pressing, but the sequence simply results in one person being created, who will live to 90. Thus, if this sequence of actions is worth taking, then a single action that results in the sequence is similarly worth taking. Consequently, one has some reason to create a happy person who will live to 90.

2.2 *Sequential desirability defended*

Sequential Desirability claims “if one should press a sequence of buttons in some order, then if another button has the same effects as pressing the buttons in that sequence in that order, then that other button should be pressed.” It is hard to deny. If two things are each worth doing, and one action has the same effects as doing both of them, then that action is clearly worth doing. This closely resembles the principle defended by Huemer (2009, pp. 466–67), who argues that whether some action is permissible cannot depend on whether it counts as one act or two acts.⁶ If Huemer’s principle is right, then if some sequence of buttons is worth pressing, so too would pressing a single button that has the same results as pressing the sequence of buttons – after all, its deontic status could not change on account of it being done as one action rather than multiple.

Denying Sequential Desirability has extremely implausible results. To see this, on such an account, there are some buttons one, two, and three, such that one should press button one and button two, but should not press button three, even though button three has the same effects as pressing buttons one and then two. Thus, one should press button one and then two, but if button three had the effect of pressing buttons one and two, by lowering a paperclip that press button one first and then button two, it would not be worth pressing. But this is absurd! If one should press both of two buttons, and there is some other action that allows them to cut out the middle man and press them in one fell swoop, that action would be similarly worth taking, all else equal. Whether pressing the buttons is wrong cannot depend on whether it is done more efficiently as one action rather than two!

In addition, denying Sequential Desirability does not avoid the sequential version of the puzzle. Defenders of the asymmetry should accept:

Sequential Asymmetry: One has no reason to take a sequence of actions that ultimately result in creating one happy person and nothing else.

After all, if one has no reason to create a happy person, then one has no reason to take a sequence of actions that just creates a happy person. However, if one accepts Generative

⁶It also closely resembles the Dominance Addition principle of MacAskill (2022, p. 181).

Improvement and Modification Improvement then they must reject this principle. After all, Generative Improvement and Modification Improvement entail that one should take a sequence of actions that culminates in creating one well-off person and has no other effects.

Finally, one might worry that this principle neglects the crucial role of intent. If the effects of an action are just what it does, then this seems disagreeable to anyone who thinks that intent affects the rightness of an act. But here the effect of an act is used in the broadest sense to include intent. So, for instance, if taking an ultimately harmful action based on good intent is right and taking another harmful action based on good intent is right again, then this would only imply that a single action motivated by both of the good intents that causes both harms would be right – which seems plausible enough. Thus, the intent worry is easily put to rest.

Even aside from this, the asymmetry is plausibly axiological in nature, rather than deontic, claiming that it is not a good thing when happy people are created. If one accepts it only deontically, then they run into the strange result that it is good when a happy person gets created but this gives one no reason to create a happy person. Normally, however, the fact that an action makes the world better gives one a reason to do it. While one could in theory reject this, this certainly seems like a cost to such a view. Thus, the argument can be reconstructed making reference solely to axiology.⁷ The principles of the argument would thus be the following:

Generative Improvement Axiology: If a button being pressed will increase the utility of an existing person by some amount and create some person with positive utility, then that button being pressed would be fortunate.

Modification Improvement: If after the button described in the Generative Improvement principle has been pressed, another button is pressed which rescinds the utility increase for the existing person, while increasing the utility of the guaranteed future person created by the first button by a much larger amount, that would be fortunate.

Sequential Desirability Axiology: If there is a sequence of buttons, each of which would be fortunate if pressed, then if there was another button that has the same effects as pressing each of the buttons in that sequence, that button being pressed would be fortunate.

Notably, these principles are almost identical to the original principles. The only difference is that these principles are axiological in nature – they are about whether it would be fortunate for such a button to be pressed rather than whether any particular person should press such a button. In fact, as long as the better-than relation is transitive, Sequential Desirability Axiology is not even needed, for then the world where the first button is pressed is better than the world where neither button is pressed, and the world where the second button is pressed also is better than the world where the first button alone is pressed, so by transitivity, the world where both buttons are pressed would be better than the world where neither button is pressed. But the only difference between the world where neither button is pressed and the world where both are pressed is that the second contains an extra happy person – thus, the addition of a happy person improves a world.

⁷ A similar argument has been given by Broome (2005).

2.3 Generative improvement defended

Generative Improvement states “If one can press a button that will increase the utility of an existing person by some amount and create some person with positive utility, then they should do so, if they will later have the option described in Modification Improvement.” This is very plausible and closely resembles the Benign Addition Principle of Huemer (2008, p. 902). If some act leaves no one worse off and some better off, it seems that it is worth taking. But pressing the button described does leave no one worse off and someone better off. The new person who exists is not worse off, because they now have a good life. The existing person whose life is improved is clearly made better off.

One might object by claiming that an action that creates a new happy person is neither good nor bad. Thus, because its value is not comparable with the value of inaction, slightly improving it does not make it more valuable than inaction. Just as if one is deciding between two incomparable spouses, the tie would not be broken by very slightly improving one of them, so too would the slightly improved option of procreation remain incomparable with failing to procreate (see, for example, Chang 1998 and De Sousa 1974). This argument, however, is quite vulnerable to a range of objections.

First, whether one’s reasons to take acts can be incomparable, as is required for this view, is deeply controversial. Many have presented arguments for why rational preferences must be complete (see Dorr et al. 2023; Gustafsson 2022; Hare 2013).⁸

Second, even if incomparability in reasons for acting exists, it is hard to believe that some actions are neither good nor bad. There might be two actions that are both very good and incomparable, or very bad and incomparable, but it is hard to imagine that some actions’ effects are neither good nor bad. Such a result would invalidate the following principle from Dorr et al. (2023, p. 438):

Strong Monotonicity: x is F . y is not F . So, x is more F than y .

This inference seems intuitively valid. If one person is tall and another isn’t, it seems the first is taller than the second. If one person is angry and another isn’t, the first is angrier than the second. But on this account, Strong Monotonicity is sometimes false. To see this, suppose we’re comparing an action that slightly benefits an existing person to one that simply creates a future happy person. Assume that whether one should take an action is determined by how worthwhile the action is. The first action would be worthwhile and the second would not be. However, the first wouldn’t be more worthwhile than the second – they’d be incomparable. Such flagrant violations of Strong Monotonicity are hard to stomach.

Third, this alone is not enough to rescue the procreation asymmetry. Even if we accept this line of reasoning, we should accept:

Extreme Generative Improvement: There exists some amount of utility, such that if one can press a button that will increase the utility of an existing person by that amount and create some person with positive utility, then they should do so, if they will later have the option to press the button described in Modification Improvement.

⁸Hare presents a paradox for incomparability, and then explores the two possible solutions to it, in chapter 3. But the paradox shows, I think convincingly, that views which affirm incomparability are quite implausible.

This principle is extremely plausible. If, for instance, creating a happy person would increase the well-being of an existing person by an extreme amount, such that it would enable them to accrue more welfare each second than the sum of all welfare experienced in the history of the world, it would be worth doing. But as long as there is any amount of benefit to existing people that would be worthwhile in exchange for creating a new happy person, the asymmetry is false. Even if the benefit must be quite large, as long as it would be worth redirecting the benefit to the newly created person if it would become larger, then one should take a sequence of actions that does nothing other than create a happy person. As a consequence, as long as the very modest Extreme Generative Improvement principle is true, combined with Modification Improvement and Sequential Desirability, the asymmetry is false. Remember, the objection to the procreation asymmetry proceeds by beginning by noting that one should benefit an existing person and create a new person as one act, and then should sacrifice the benefit to the existing people to provide a larger one to the guaranteed future person. For this to work, it need not be that any benefit to an existing person is worth creating a happy person but merely that some benefit to an existing person can be worth creating a happy person.

Fourth, if we simply consult our intuitions about such a case, it seems overwhelmingly plausible that it is good to press the button described in the scenario. Pressing the button leaves no one worse off, one person better off, and creates a happy person. It follows from a slightly modified version of the Modal Pareto principle given by Huemer (2008, p. 468):

Deontic Modal Pareto: If an action benefits some people and harms no one, such that no one's life is worse after the action was taken than it would have otherwise been, and some people's lives are better, the action should be taken.

Such a principle is extremely plausible. And yet it straightforwardly entails that one should press the first button; such an action improves the life of an existing person while making no one's life worse. In addition, the case described is completely different from other cases where actions are incomparable. In other cases, there are lots of strong reasons pointing in both directions, and the magnitudes of the reasons are said to be incomparable. But in this case, it is not clear that there is any reason not to take the act, while there is at least one good reason to take the act, for no one is wronged by the creation of a happy person.

One might object to this claim and say that, given that the person will experience lots of suffering over the course of their life, there are strong reasons not to take the act. The reasons not to create someone who will suffer are incomparable with the reasons to create someone who will experience a large amount of welfare. Thus, in a case where one will be created who will suffer a lot but experience greater welfare, such an action is neither good nor bad. However, in the aforementioned scenario, the person need not suffer. We can stipulate that this button creates a person with some positive welfare and no suffering. Such an action is not bad – even Benatar's (2006), famous procreation asymmetry would not hold that doing this is wrong. There is no asymmetry strong enough to hold that creating a happy person who never suffers is bad. But that is the only assumption needed for the argument.

One might think that, while ordinarily one should press the button described in Generative Improvement, this is only so if they will not have the option to press the extra button described in Modification Improvement. Creating a happy person and

benefitting someone is ordinarily worth doing, but if after one does it, they will then have a reason to rescind the benefit, then they know that what the sequence of acts will result in will not be worthwhile, because it will just create a happy person. As such, they have no reason to press the first button if the second button will be available later. But the claim that the button described in Generative Improvement is only worth pressing if the button described in Modification improvement will not be able to be pressed is not plausible. Imagine the following dialogue:

Person 1: Hey, I think I will create a person by pressing a button. It is no strain on us and it will make the life of a random stranger better, so it will be an improvement for everyone.

Person 2: Oh, great! I will offer you a deal. If you do that, and you do not give the gift to the stranger, instead, I'll give your baby 20 million times as much benefit as the gift would have.

Person 1: Thanks for the offer. It is a good offer, and I would take it if I were having the baby. But I am not having a baby now, because you offered it.

Person 2: What? But you don't have to take the offer.

Person 1: No, no. It is a good offer. That is why I would have taken it. But now that you have offered me this good offer, that would be worth taking, it makes it so that having a baby by button is not worth doing.

Clearly, person 1 is being irrational here (for a similar principle, see Hare 2016, p. 460). If, after taking some action, one gets another good option, that would not make the original action not worth taking. The fact that some action allows one to do other worthwhile things counts in favor of it, not against it. As Huemer (2013, p. 334) notes, it is perfectly rational to refuse to take an action because you predict that if you take it, you will do other things that you should not. But it is clearly irrational to refuse to take an action on the grounds that, if you do, you will do other worthwhile things.

A person might appeal to backward induction to try to avoid this argument. Backward induction involves one maximizing their payouts at each choice node while taking as given what they will do at future choice nodes. In this case, one would refuse to press the first button, because they know that if they do so they would subsequently press the second button, and pressing both buttons simply creates a happy person which does not advance their interests. While this is perhaps the most promising way to reject the argument, it has major costs.

First, it involves giving up on Generative Improvement. On this picture, one should not press the first button, even though pressing it and stopping would be conducive to one's goals, because it will equip them with a later option to press another button that will, at the time, be worth pressing. But, as previously discussed, it is hard to believe that one should not take an action simply because if they take the action it will give them another action that is worth taking.

Second, while one's evaluations might change over time, one who rejects Generative Improvement as a way of avoiding the conclusion must think that prior to pressing the button described in Generative Improvement, one should oppose pressing the second button in Modification Improvement, and then they should change their mind when they get to Modification Improvement. Prior to pressing the first button, one is indifferent between pressing both buttons and pressing no buttons, and prefers pressing just the first button to pressing no buttons. Therefore, if one holds on to the idea that the better than relation is transitive, they must think that the agent's normative evaluation

should change after pressing the first button (they must, after all, think that though the agent prefers pressing the second button and the first to just pressing the first, that the agent takes the fact that pressing the first would enable them to press the second button to count against it).

But this is implausible in two ways. First, it implies that prior to taking the action described in Generative Improvement, one should hope they act wrongly when given the option in Generative Improvement. But plausibly, one should not hope they do, in the future, something that they should not do. Second, it implies counterintuitively that one should, prior to pressing the button described in Generative Improvement, prefer a world where only that button is pressed to the intuitively superior world where they also press the Modification Improvement button.

Third, given that one will, on this picture, either be indifferent between pressing no buttons and pressing both buttons, such a view violates the version of transitivity articulated by Gustafsson (2022), according to which if $A > B \sim C$ then $A > C$ (\sim denotes incomparability). This makes it vulnerable to the money-pump of Gustafsson (2022, pp. 39–44).

Fourth, such a view just seems irrational independently. The right principle of rationality should not result in one losing out on what is valuable when alternative sequences of actions could result in them gaining those great goods. But this view does, resulting in one losing out on pressing the first button, which would be valuable, simply because it will subsequently enable them to press the second button. In an even more troubling example of this, consider a case where the first button would improve the welfare of an existing person by 10 units of well-being and create a happy person with 10 units of well-being, as well as enabling the pressing of the second button. Suppose that the person is given the option to press another button which would improve the welfare of an existing person by 8 units of well-being and would increase the well-being of an existing person by 8 units, but would not enable the pressing of the second button. On such a view, one should press the inferior button, that produces only 8 units of benefit to existing people, even though simply pressing the first button is strictly superior to this.

2.4 Modification improvement

The final premise in the argument, Modification Improvement, states “if one can, after pressing the button described in the Generative Improvement principle, press another button, which will rescind the utility increase for the existing person, while increasing the utility of the guaranteed future person created by the first button by an amount that is many times larger, they should do so.” This premise is probably the least controversial of the premises. It follows from:

Superior Benefit Offspring: If one is choosing between providing some benefit for a stranger or a more significant benefit for their own offspring who is guaranteed to exist in the future, they should provide the greater benefit for their offspring.

Many claim that we have special obligations to our offspring. And yet even those who deny this think that if we can benefit either our offspring or a stranger, where the benefit to our offspring would be greater, we should produce the greater benefit for our offspring. If we think back to the case given in section 2.1, where one button can create a happy person and enable someone to live an extra year, and another button can

rescind the increase in life, while allowing the offspring to live an extra 20-years, it seems clearly immoral not to press the second button. Once the person is guaranteed to exist, benefitting them is clearly more important than producing much smaller benefits for disparate strangers.

One could deny such a claim. It is not clear that there is an independent motivation for this, but one who is sufficiently committed to the procreation asymmetry could try to reject this principle to salvage their view. However, even if Modification Improvement is false, the following more modest principle is virtually undeniable:

Extreme Modification Improvement: There is some amount of benefit such that if one can, after pressing the button described in the Generative Improvement principle, press another button which will rescind the utility increase for the existing person, while increasing the utility of the guaranteed future person created by the first button by that amount, they should do so.

This principle merely holds that rescinding the benefit provided by the button described in Generative Improvement is worth doing if it provides some sufficiently large benefit to the guaranteed future person. Such a premise can be deduced from:

Extreme Superior Benefit Offspring: If one is choosing between providing some benefit for a stranger or a benefit that is a million times greater for their own offspring who is guaranteed to exist in the future, they should provide the benefit for their offspring.

One could object by accepting Extreme Modification Improvement and Extreme Generative Improvement. However, they might adopt a capped view of well-being, according to which there is a limit to how well-off one can be.⁹ Thus, if to justify procreation, one has to produce a very significant benefit to existing people, and one should only redirect the benefit from existing people to the guaranteed offspring if the benefit is much greater, it might be that, because there is a limit to how well off one can be, there is in fact no benefit large enough for vetoing the benefits to be worth it for the sake of the guaranteed future person.

To illustrate this, suppose that the maximum amount of well-being one can have is 100,000 units. Suppose that to justify creating a person who would have 10,000 units of well-being, one would have to produce 50,000 units of well-being for existing people. Now suppose that to justify redirecting a benefit from an existing person to a guaranteed future person, it would have to be four times as great. On such a view, there is no sequence of actions that is worth taking that only results in creating a very well-off person, because for the principles to be justified, the person who is created would have to be better off than anyone could be.

Notably, such a conclusion requires a whole host of controversial assumptions. For such a solution to work, it would require accepting a cap on the value of well-being, denying Deontic Modal Pareto, denying Superior Benefit Offspring, and denying Generative Improvement. All of these assumptions with the possible exception of accepting a cap on the value of well-being are hard to stomach.

This result becomes even more implausible in concrete cases. For instance, suppose that there are two buttons, the first of which creates a supremely happy person who is

⁹See Temkin (2012, sections 10.6–10.7) and Williams (2022).

never miserable and also enables an existing person to live an additional 70 years of happy life. Such a button would clearly be worth pressing. But if there was another button that rescinded the 70 years of well-off additional life (before it was received, of course) while enabling the guaranteed future offspring to live an extra 300 years of supremely well-off life, that would clearly be worth pressing.

3. From the sequence argument to strong reasons to procreate

3.1 *Where are we now?*

So far, we have established from extremely minimal axioms that one has at least some reason to create a happy person. In fact, we can deduce that one has a reason to take a sequence of actions that creates a slightly happy person and does nothing else from Deontic Modal Pareto which states “if an action benefits some people and harms no one, such that no one’s life is worse after the action was taken than it would have otherwise been, and some people’s lives are better, the action is worth taking,” and Superior Benefit Offspring which states “if one is choosing between providing some benefit for a stranger or a more significant benefit for their own offspring who is guaranteed to exist in the future, they should provide the greater benefit for their offspring.” If these are combined with Sequential Desirability, we can deduce that if there was an action that would simply create a minimally happy person and do nothing else, it would be worthwhile.

For instance, suppose we want to show that one should create a person with one microutil worth of well-being (that is a very small amount, which we will assume is equal to the amount of pleasure one gets from eating chocolate for ten minutes). It is worth pressing a button that gives an existing person half a microutil and creates a person who will have 0.1 microutils worth of happiness. Then, it is worth pressing another button that eliminates the half a microutil of benefit that would be given to existing people in order to increase the happiness of the now guaranteed-to-exist person by 0.9 microutils. Therefore, it is worth pressing a sequence of buttons that creates a person with a microutil worth of happiness, and from Sequential Desirability, we can derive that one has some reason to create a person with a microutil worth of happiness.

Even if one rejects Deontic Modal Pareto and Superior Benefit Offspring, nearly irresistible principles can be employed, which together entail that one has a reason to create a very happy person. Such a result is extremely significant. The procreation asymmetry thus conflicts with a series of extremely widely held principles. Consequently, arguments for it, such as the alleged impossibility of acting for the sake of a nonexistent person must be wrong.¹⁰

But suppose that one accepts Deontic Modal Pareto, Superior Benefit Offspring, and Sequential Desirability, each of which are extremely plausible, yet not quite as irresistible as their more modest versions. From these, we can go further and derive the existence of very strong reasons to procreate and that one has some reason to create any person whose life will contain positive aggregate utility.

3.2 *Do not just make people happy, make net happy people*

The argument so far has proceeded by starting with plausible premises that entail the falsity of the procreation asymmetry. Then, I have shown that even if we weaken

¹⁰For a discussion of this argument, see, for example, Frick (2020, p. 56).

those premises to make them almost undeniable, they still are enough to disprove the procreation asymmetry. The argument in this section will assume the stronger versions of these premises, which are themselves still quite weak, and show how from those we can establish more significant results.

So far, we have established that one should create any person whose life contains positive well-being and no suffering. But one might think that this is not such a significant result – over the course of one's life, even very well-off people will suffer a lot. Maybe this means procreation is fine in utopia, but it does not vindicate procreation in normal cases, where lives will contain lots of both well-being and suffering.

But we can go further from these assumptions. We have already seen reasons to support the Deontic Modal Pareto principle. But from that principle, as well as the others argued for, we can derive that one has a reason to bring into existence anyone with a net positive life. The reason to think this can be seen in four steps:

- 1) One should press a button that creates a person with one unit of well-being and increases the well-being of an existing person by one unit. Assume one unit of well-being is not very much, perhaps the amount of well-being one gets, on average, from a nice vacation. This follows from Deontic Modal Pareto.
- 2) One should press another button that eliminates the benefit to the existing person while increasing the well-being of the guaranteed future person by two units. This follows from Superior Benefit Offspring.
- 3) One should press a third button which increases the well-being of the guaranteed future person by 46 units while increasing their suffering by forty-five units. This follows from Deontic Modal Pareto; assuming one unit of well-being benefits someone as much as one unit of suffering harms them – which we can stipulate by defining the units that way – this would leave the newly created person better off and harm no one.
- 4) Therefore, one should press a sequence of buttons which creates a person with 48 units of well-being and forty-five units of suffering, they should do so.
- 5) Therefore, by Sequential Desirability, one should press a single button that creates a person with forty-eight units of well-being and forty-five units of suffering.

It is worth noting that this argument is fully general. By changing the numbers around, this broad argumentative schema can be used to derive that one should press a button that would create anyone who has positive aggregate welfare, no matter how much they suffer, as long as their well-being is greater, assuming such an action will have no other effects.

3.3 It is good to create people even at significant personal cost if the child's welfare is greater than the cost

Suppose that one is in a position to create a person at a significant personal cost. Suppose that the harm to the person who has the child would, while significant, be less great than the well-being level of the child. Most people think that such an action would not be worth taking. Not only would it not be obligatory, it would not be a worthwhile action – even if one has some reason to create happy people, their reasons cannot be strong enough to justify very significant personal costs. Thus, many would regard this action as equivalent to giving up one's own life to save the leg of someone else. But the argument I have given so far gives us a good reason to reject this.

The previous section has established that one has reason, *ceteris paribus*, to create any person whose life will have positive well-being overall. Thus, one has a reason to press a button that will create a person whose life contains significant suffering, as long as it contains greater happiness. The following principle is additionally plausible:

Offspring Agony Passing On: one should endure some amount of suffering as long as it averts a greater amount of suffering from being experienced by their offspring.

Notably, this premise does not claim that one is required or obligated to pass on such agony. The should used here is the should of most reason (see Chappell 2020). For those who dispute that this is a legitimate use of the term should, consider replacing every use of the term should with “has most reason to do,” and the argument works just as well. Thus, this premise does not raise worries about demandingness – it is perfectly compatible with this premise that such an action is supererogatory but not obligatory.

With this in mind, such a premise is very plausible. Enduring some cost to prevent another from bearing a greater cost is a good action – it is an action that one has more reason to take than not to. This is especially plausible when the person whose suffering one can avert is one who they have created! But if this is true, then one should (in the sense of most reason) endure some amount of suffering to create anyone with aggregate positive welfare. To see this, consider the following sequence of actions:

- 1) A person can press a button that creates a person with 50 units of well-being and 49.9 units of suffering.¹¹ This is worth doing, as section 3.2 showed.
- 2) The person can then press a button that would relieve 40 of those units of suffering but would cause themselves to endure 39 units of suffering. This is worth doing, as shown by Offspring Agony Passing.
- 3) Therefore, one has a reason to press a sequence of buttons that creates a person with 50 units of well-being, 9.9 units of suffering, and would cause the button presser to endure 40 units of suffering.
- 4) Therefore, by Sequential Desirability, one has a reason to press a single button that creates a person with 50 units of well-being, 9.9 units of suffering, and that causes them to endure 40 units of suffering.

This conclusion generalizes. Because it is always worth creating a person with net aggregate welfare, and it is worth passing suffering from one’s offspring to oneself as long as the amount of suffering is diminished at all, one has reason to create any person who will experience more positive welfare than the creator goes through in creating them.

Even if one rejects the claim that one should create any person with positive welfare, the conclusion can still work, albeit in a more modest form. The more modest version of the argument against the asymmetry showed that a person should create one who is very happy. But if this and Offspring Agony Passing On are both true, then one has a reason to create someone if the difference between their level of well-being and the level of well-being that makes a person worth creating is greater than the costs of creating them to the one who creates them. Such a result is still quite significant.

¹¹Remember one unit of well-being is n’t very much, perhaps the amount of well-being one gets, on average, from a nice vacation.

3.4 *The Repugnant Conclusion*

One might notice a certain similarity between this argument and the Benign Addition argument given by Huemer (2008, p. 902) based off one from Parfit (1984, Ch. 19). This principle is used to support the Repugnant Conclusion which says that a world with 10 billion people living great lives is worse than one with a sufficiently huge number of people living lives barely worth living. One might suspect this entails the Repugnant Conclusion in the following way.

The argument in section 3.3 establishes that it is good to create a happy person if the created person's well-being level is greater than the amount of well-being sacrificed in creating them. But this implies that if after 10 billion well off people get created, they have a slight dip in their well-being level to create a barely happy person, this improves the world. Even if it is only worth creating a person if the cost to do so, in terms of well-being, is 1% as great as their well-being level (for example, the world is improved by creating a person with 100 microutils if and only if that costs less than 1 microutil to existing people), by iterating the process, eventually, one gets the Repugnant Conclusion.

In response, one has two ways to go. First of all, they can simply accept the Repugnant Conclusion. Lots of people already do this because it is implied by lots of different plausible arguments.¹² Given that lots of people accept the Repugnant Conclusion for lots of other reasons, if the minimal principles described here also imply it, that just seems to be further reason to accept it.

Alternatively, one could reject Modification Improvement or Generative Improvement. They could hold that one should not always create an existing person with a positive welfare level even if it benefits an existing person. Perhaps one has some moral reason not to create a just barely well-off person. This would not rescue the asymmetry, however, because my argument against the asymmetry requires the much more moderate Extreme Generative Improvement and Extreme Modification Improvement.

4. Conclusion

Common-sense morality seems to claim that one has some reason to make existing people happy but not to create happy people. This judgment is known as the procreation asymmetry and is controversial; various people have struggled to find a comprehensive defense of it and it has been extensively criticized. This paper shows that the asymmetry is indefensible because it is committed to the view that one should take a sequence of actions that it holds one has no reason to take.

Every plausible view will hold that benefiting guaranteed future people will be worthwhile and that creating a happy person while benefitting an existing person will be similarly worthwhile. But once one has created a happy person and benefitted an existing person, there is a future person who is guaranteed to be created; such a person should be benefitted, even at the cost of rescinding the benefits to the existing people. But in combination, these judgments produce the result that one should take a sequence of actions that simply creates a happy person. This in combination with modest principles runs sharply counter to the asymmetry. The asymmetry has thus sowed the seeds for its own demise by instructing one to take a sequence of actions that it claims one has no reason to take.

¹²See, for example, Huemer (2008), Tannsjo (2002), and Spears & Budolfson (2021).

The argument in this paper, once in place, enables us to derive even more extreme results; not only does one have some reason to create a happy person, they have a very strong reason to create one who would live a good life. From these minimal axioms, we can see that one has a reason to create any person who has positive aggregate well-being and that creating one is worthwhile as long as the costs to the creator are less than the aggregate well-being that the created would get over the course of their life.

References

- Beckstead, Nick** (2013). 'On the Overwhelming Importance of Shaping the Far Future.' Ph.D. thesis. New Jersey: Rutgers University.
- Benatar, David** (2006). *Better Never to Have Been: The Harm of Coming into Existence*. Oxford: Oxford University Press.
- Broome, John** (2005). 'Should We Value Population?', *The Journal of Political Philosophy*, **13**(4), pp. 399–413.
- Chang, Ruth** (1998). *Incommensurability, Incomparability, and Practical Reason*. Harvard: Harvard University Press.
- Chappell, Richard Yetter** (2017). 'Rethinking the Asymmetry', *Canadian Journal of Philosophy*, **47**(2–3), pp. 167–77.
- Chappell, Richard Yetter** (2020). 'Deontic Pluralism and the Right Amount of Good.' In *The Oxford Handbook of Consequentialism* ed. Douglas W. Portmore. pp. 498–512. Oxford: Oxford University Press.
- De Sousa, Ronald** (1974). 'The Good and the True', *Mind*, **83**(332), pp. 534–51.
- Dorr, Cian, Nebel, Jacob, and Zuehl, Jake** (2023). 'The Case for Comparability', *Noûs*, **57**(2), pp. 414–53.
- Frick, Johann** (2020). 'Conditional Reasons and the Procreation Asymmetry', *Philosophical Perspectives*, **34**(1), pp. 53–87.
- Grill, Kalle** (2017). 'Asymmetric Population Axiology: Deliberative Neutrality Delivered', *Philosophical Studies*, **174**(1), pp. 219–36.
- Gustafsson, Johan** (2022). *Money-Pump Arguments*. Cambridge: Cambridge University Press.
- Hare, Casper** (2013). *The Limits of Kindness*. Oxford: Oxford University Press.
- Hare, Casper** (2016). 'Should We Wish Well to All?', *Philosophical Review*, **125**(4), pp. 451–72.
- Huemer, Michael** (2008). 'In Defence of Repugnance', *Mind*, **117**(468), pp. 899–933.
- Huemer, Michael** (2009). 'A Paradox for Weak Deontology', *Utilitas*, **21**(4), pp. 464–77.
- Huemer, Michael** (2013). 'Transitivity, comparative value, and the methods of ethics', *Ethics* **123**(2), pp. 318–45.
- MacAskill, William** (2022). *What We Owe the Future*. Oxford: Basic Books.
- Meacham, Christopher** (2012). 'Person-Affecting Views and Saturating Counterpart Relations', *Philosophical Studies*, **158**(2), pp. 257–87.
- Narveson, Jan** (1973). 'Moral Problems of Population', *The Monist*, **57**(1), pp. 62–86.
- Parfit, Derek** (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Pummer, Theron** (2016). 'Adding Happy People.' In *Philosophers Take on the World*, ed. David Edmonds. pp. 236–39. Oxford: Oxford University Press.
- Spears, Dean, and Budolfson, Mark** (2021). 'Repugnant Conclusions', *Social Choice and Welfare*, **57**, pp. 567–88.
- Tannsjö, Torbjörn** (2002). 'Why We Ought to Accept the Repugnant Conclusion', *Utilitas*, **14**(3), pp. 339–59.
- Temkin, Larry Temkin** (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press.
- Williams, Bernard** (2022). 'The Makropulos Case: Reflections on the Tedium of Immortality.' In *The Metaphysics of Death*, pp. 71–92. Stanford: Stanford University Press.
- Wolf, Clark** (1996). 'Social Choice and Normative Population Theory: A Person Affecting Solution to Parfit's Mere Addition Paradox', *Philosophical Studies*, **81**(2), 263–82.