

This is the penultimate draft of a paper forthcoming in *Philosophy and Phenomenological Research*. Please don't cite this version.

PESSIMISM AND PROCREATION¹

DANIEL PALLIES
Lingnan University

The pessimistic hypothesis is the hypothesis that life is bad for us, in the sense that we are worse off for having come into existence. Suppose this hypothesis turns out to be correct — existence turns out to be more of a burden than a gift. A natural next thought is that we should stop having children. But I contend that this is a mistake; procreation would often be permissible even if the pessimistic hypothesis turned out to be correct. Roughly, this is because we are often in a position to know that future people will approve of having been created, and their approval will not be inappropriate even if they are worse off for having been created. Thus, our respect for the attitudes of future people can permit us to create them.

1. Introduction

Most theories of well-being enumerate both goods and bads: things which increase our well-being (or make us better off), and things which decrease our well-being (or make us worse off). For example, hedonism tells us that pleasant experiences increase our well-being, and unpleasant experiences decrease our well-being. Theories with this structure should admit the possibility that overall, our lives include more bad than good — more displeasure than pleasure, for example. And if this possibility obtains, then it is natural to suppose that life is *bad for us*, in the sense that we are *worse off* for having come into existence. Existence would turn out to be more a burden than a gift. Call this the *pessimistic hypothesis*.

Supposing that the pessimistic hypothesis is correct, a natural next thought is that we should stop having children. This natural thought has been expressed in different ways by a number of philosophers. David Benatar argues that we are always harmed and never benefitted by coming into existence, and takes this to imply that it is impermissible to create new people (2006). Similarly, Derek Parfit assumes that it sometimes *is* permissible to create new people, and takes this to imply that some people's lives are good for them (1984).² But I contend that both Benatar and Parfit are making a mistake — it would often be permissible to create new people, even if the pessimistic hypothesis turned out to be correct. Roughly, this is because many people will *approve* of having been created, even if their creation made them worse off, and our respect for their attitudes towards their creation can permit us to create them. In other words, insofar as

¹ This paper was presented at the 2022 meeting of the Pacific Philosophical Association as well as the 2023 Ethics Forum at California State University Long Beach. Thanks very much to Kirsten Egerstrom, Molly Gardner, and audiences at those talks for their insightful and constructive comments. Thanks also to Mark Schroeder, Ralph Wedgwood, John Hawthorne, an anonymous reviewer at PPR, and — as always — Jennifer Foster for extended discussion on earlier drafts of this paper.

² Parfit holds that, if some future person would have a life that is “not worth living”, then this would be a strong reason *not* to create them. The meaning of the phrase “not worth living” is admittedly not entirely clear. But according to one common usage of Parfit's, if one is worse off for having been born, then one's life is “not worth living” (p.358-359, 391).

future people will approve of having been created, we can to that extent ignore the effects of their creation on their well-being.

I begin in §2 by explaining how I understand the pessimistic hypothesis and why it is worth taking seriously. I then turn to my argument, which has two parts. First, in §3, I defend the general principle that it can be permissible to act in ways that make people worse off because we are justified in believing that those who are made worse off will approve of our having acted as we did. Second, in §4, I argue that the general principle applies to the case of procreation. Many parents can be justified in believing that the people they create will approve of having been created, even if the pessimistic hypothesis is true. For this reason, many parents' respect for future peoples' attitudes can make it permissible for them to create those people. In §5 I contrast my arguments with one raised and criticized by Seana Shiffrin (1999) and Asheel Singh (2018). Having shown that their criticisms do not threaten my argument, I briefly summarize my conclusions in §6, while also sketching some implications for the so-called "procreation asymmetry".

2. The Pessimistic Hypothesis

The pessimistic hypothesis is a hypothesis about our well-being, where *well-being* is the kind of value at issue when we say that someone's life is going well or badly for them. Well-being, so-understood, bears a conceptual connection to *egoism*: insofar as our motivations are *egoistic*, all we care about is maximizing our own well-being. Thus, different theories of well-being can be understood as issuing different recommendations to egoists. Egoistic *hedonists* will attempt to maximize their pleasure and minimize their pains. Egoistic *objective list theorists* will attempt to maximize their share of objective goods: friendship, achievement, knowledge, etc. Both groups want to be as well-off as possible, but they have different ideas about how this can be accomplished.

The pessimistic hypothesis is the hypothesis that, from an egoistic perspective, coming into existence was a bad thing:

Pessimistic Hypothesis. It would be better *for us* if we did not exist.

It would have been better for us (for *all* of us, or at least *most* of us) if we had never been born. My goal is not to defend this hypothesis, but to defend a *conditional* claim about the implications of this hypothesis, supposing that it is true. I claim that, even if pessimism is true, it is often permissible to create new people. This should not be interpreted as the claim that, *if things were otherwise* — if people experienced greater degrees of suffering and failure, for example — then it would still be permissible to create new people. The claim is that, *if the pessimistic hypothesis turned out to actually be true*, then it will still be permissible to create new people. Borrowing terminology from two-dimensionalist semantics, the goal is to consider the implications of the pessimistic hypothesis "considered as actual" — as something that might *turn out to be true* — not "considered as counterfactual" — as something that *could have been true* (Chalmers 2002).

Strictly speaking, the truth of the pessimistic hypothesis is irrelevant to my defense of this conditional claim. But the interest of the conditional claim obviously depends on the plausibility of pessimism. If pessimism is utterly implausible, then the conditional claim is uninteresting. But pessimism is not utterly implausible, and the rest of this paper can be read as an argument to that effect. Pessimism seems implausible at least in part because it seems to have certain deeply implausible consequences — for example, that no one should have children. In the rest of this paper I will argue that pessimism does not

have this and other related consequences; to the extent that I am successful, this should make pessimism seem less implausible.

Furthermore, pessimism can also be defended in a more direct way, by thinking through some actual theories of well-being. Consider hedonism, for example. In its simplest and best-known form, hedonism says that your life is good for you to the extent that it is pleasant, and bad for you to the extent that it is unpleasant. If this version of hedonism is true, and the pessimistic hypothesis is true, then that means our lives are more unpleasant than pleasant. This would not entail that our lives are filled to the brim with agony. It might be that the dull discomforts of hunger outweigh the pleasures of eating; the displeasures of uncomfortable temperatures outweigh whatever pleasures we get from being at the perfect temperature; and so on. This version of pessimism is not obviously correct, but neither is it obviously false. Speaking for myself, I do not have a strong sense of the hedonic balance of my life. There is, after all, quite a lot that goes into the balance: every little affective twinge of every moment of every day. If we're at all modest about our powers of introspection — as we should be!³ — then we should at least be open to the possibility that our hedonic balance comes out negative. So, if we are open to the possibility that hedonism is the correct theory of well-being, we should be open to the possibility of pessimism.

Similar considerations apply to the desire satisfaction theory. If pessimism and the desire satisfaction theory are correct, then the disvalue of desire frustration outweighs the value of desire satisfaction. That hardly seems like an obvious falsehood. After all, it is far from obvious that our desires are more often satisfied than frustrated. It would be rather surprising if it turned out that *usually*, when we desire something, we get it.⁴ So pessimism should be regarded as a serious possibility on standard versions of both hedonism and desire satisfactionism, two of the leading theories of well-being.

There are other theories of well-being which might seem to make pessimism a less serious possibility. There are various versions of the objective list theory and other pluralist theories, which tell us that many different kinds of things are good for us — not just pleasure or happiness but also knowledge, achievement, friendship, meaningfulness, life satisfaction, etc. Compared to hedonism and desire satisfactionism, these pluralistic theories provide a more expansive picture of all that is good in life, and it is easy to get the impression that this more expansive picture is also more *optimistic* about the goodness of life, in the sense that if it is accurate then life is likely good for us. But we should not neglect the “negative side” of pluralistic theories. To the extent that the putative goods are *prima facie* good, their opposites are *prima facie* bad — ignorance, failure, conflict, meaninglessness, life dissatisfaction, etc. And it is far from obvious that our lives contain more of the putative goods than their opposites.⁵ Common sense tells us that we fail more often than we succeed, our ignorance is more expansive than our knowledge, etc. Furthermore, insofar as

³ For philosophical arguments in favor of modesty, see Schwitzgebel 2006, and for those arguments applied to the case of pleasure and pain specifically, see Bramble 2013.

⁴ Most desire satisfactionists hold that only *some* desires are relevant to well-being — candidates include self-interested desires, affective desires, or informed desires (Heathwood 2019). For these theorists, what matters is not the balance of satisfaction and frustration among *all* our desires, but only among *relevant subset* of desires. Still, the general point remains the same: for any plausible subset of “relevant” desires, it would be surprising if it turned out that those desires are usually satisfied.

⁵ In recent years a number of philosophers have made a start at identifying the opposites of pluralistic goods, and have argued that these identified opposites are indeed non-instrumentally bad for us (Kagan 2014; Campbell & Nyholm 2015; Hurka 2020). Christopher Rice demurs in arguing that at least some of the opposites (false beliefs, failed projects, and unhealthy relationships) are at most instrumentally bad for us (2019).

one is tempted to accept a pluralistic theory *because* it is optimistic, this is presumably because pessimism is seen as having implausible consequences. I will be arguing that pessimism does not have the implausible consequences that it is thought to have, so these arguments serve to at least partially undercut the motivation for adopting a more optimistic theory of well-being.

With all that said, there is an important non-theory-driven motivation for rejecting pessimism, which is simply that most people are glad to be alive. After all, there is abundant sociological evidence to the effect that most people are mostly satisfied with their lives (Veethoven 2020; Blanchflower, & Oswald 2017). Most of us, most of the time, do not lament the fact that we are alive. Doesn't this strongly suggest that our lives are not bad for us — in other words, that pessimism is false?⁶

The rest of this paper can be read as a response to this line of thought. On many theories of well-being, pessimism is consistent with the claim that life is a mixed bag of pains and pleasures, satisfaction and frustration, success and failure, etc. I will argue that if life is a mixed bag, then we should expect that most people will approve of having come into existence — and indeed that it is not inappropriate for us to approve of having come into existence.⁷ So the fact that people approve of existence is not strong evidence against pessimism, and it does not put strong pressure on us to accept a theory of well-being on which pessimism is unlikely to be true. In fact, I will ultimately suggest that pessimism may do *better* than optimism with respect to accommodating the deliverances of common sense morality, particularly the so-called “procreation asymmetry”.

The general upshot is that we have less to fear from pessimism than is ordinarily supposed. The truth of pessimism would not overturn the deliverances of common sense morality — it would not suggest that it is inappropriate to approve of having been born, nor would it suggest that it is wrong to have children. So our commitment to common sense claims about life and procreation should not lead us to reject pessimism. Pessimists can bring new people into the world with a clear conscience.

3. Step One: Well-being and Attitudes

Pessimism is a thesis about well-being, but well-being is not all that matters for the ethics of procreation. Future people's attitudes matter, too. We are often justified in believing that future people will *approve* of having come into existence, and this makes a difference to whether or not we should create them.

The general idea is as follows. In general, if your acting in a certain way would make someone worse off, then that is a reason to refrain from acting in that way. But if the person made worse-off will approve of your having acted as you did, then the fact that they will approve of your action can make it the case that you have less reason to refrain from acting in that way. As a result, having the person's future approval can make it permissible to make them worse off. (*Approval*, as I understand it, is a gradable psychological attitude — the opposite of *wishing it had been otherwise*.) Creating a new person, on the pessimistic hypothesis, is an action that makes someone (the created person) worse off. But if the created person will approve of having been created, then this can make it permissible for us to create them. Many potential

⁶ Thanks to an anonymous reviewer for raising this worry.

⁷ Furthermore, I will argue that no plausible theory of well-being can rule out the existence of cases in which it is reasonable for one to be glad that one is living one's actual life as opposed to an alternative that would be better for one (and worse for no one.) So even if one is antecedently committed to a theory of well-being on which pessimism is extremely implausible, one should admit that there can be *other* cases, outside the special case of procreation, in which it is not inappropriate to be glad that one living one's actual life as opposed to an alternative that would be better for one (and worse for no one). This, I think, is an interesting conclusion in its own right.

parents are indeed justified in believing that the people they create will approve of having been created. So it is permissible for those potential parents to create those people.

To illustrate the general point regarding the relationship between making someone worse off and their future approval, consider the following case:

Meteor Shower. You are watching a meteor shower late at night. You are alone, but you would rather watch with someone else. You could wake your dad so that the two of you can watch together, but you know that he has not been feeling well lately, and he needs his rest. He will undoubtedly be worse off if you interrupt his sleep, and the costs of being awoken in the middle of the night will be significant. His symptoms — nausea, headaches, muscle pain, whatever — will be intensified and prolonged. But you *also* have excellent reason to believe that your dad will *approve* of having been woken up, despite the fact that he will be worse off. He did not *ask* you to wake him; in fact he had no idea that there would be a meteor shower tonight. But you know him well enough to be justified in thinking that he will be glad to see the meteor shower, even if he knows that he would be better off getting a full night's rest.

One may find it tempting to think that, in the long run, your dad *will* be better off for watching the meteor shower. After all, it is highly plausible that sharing an experience of natural beauty with one's child is the sort of thing that contributes to a good life. But it is crucial that we not interpret the case as one in which your dad is better off in the long run for having been awoken to watch the meteor shower. It is crucial that your dad really will be worse-off as a result of being awoken early. To secure this interpretation, we can suppose that if he is awoken early he will miss some *other* experience of natural beauty.⁸ Perhaps you would have shared a cup of coffee and watched the sunrise, but if your dad watches the meteor shower he will sleep through the sunrise. Even so, he might still be glad you woke him to watch the meteor shower, and you might be justified in believing this.

There are any number of reasons why your dad might approve of having watched the meteor shower rather than the sunrise, even though he would be better off watching the sunrise. For example, he might approve in part because he benefited *you* by watching the meteor shower with you. Perhaps he knows that you would have felt lonely if you were to watch the meteor shower by yourself, and his concern for you leads him to approve of the fact that he watched the meteor shower with you. In that case, concern for *his own* well-being is not what leads him to approve of having been awoken in the middle of the night. But he approves nonetheless. After all, your dad is not a pure egoist. He might approve of something happening to him despite the fact that, if it had not happened, he would be better off. One reason is that he cares about your well-being in addition to his own well-being.

On the other hand, his approval might not be based on any thoughts about well-being at all. Not all our reasons come from considerations of well-being. Having watched the meteor shower with you, and having enshrined the memory of the experience, he might simply approve of that experience for its own sake. If so, then even if it occurs to him that he would have been better off getting more sleep and watching the sunrise, this will not make him stop approving of having watched the meteor shower. In a reflective mood,


⁸ More generally, for *any* goods that your dad receives by watching the meteor shower, we can imagine that his life would have included different goods of the same kind if he had missed the meteor shower. In this way we can accommodate most any theory of well-being, while preserving the crucial point that your dad is worse off for watching the meteor shower.

he might say: “The meteor shower was beautiful, and I treasure the memory of watching it with you. I know that the sunrise would have been beautiful, too, and if we’d watched it I’d treasure *those* memories just as much. But I don’t regret that you woke me up. I didn’t regret it then — despite the resulting headache and nausea — and I don’t regret it now, even knowing that we could have had a different, equally meaningful experience instead.”

Again, your dad is not a pure egoist. He sometimes approves of things that make him worse off, either because those events are good for other people, or because he cares about those events in ways that go beyond their contributions to well-being. Like most of us, his attitudes of approval are not always backed by well-being calculations — he might continue to approve of A rather than B even if it occurred to him that B would be more conducive to well-being. In particular, he might feel this way about you waking him up, rather than letting him sleep. And you might be justified in believing all this about him.

It is not hard to imagine that in a case like this, it could be permissible for you to wake your dad. There are some details to fill in, and as we will see the details matter for whether or not it is permissible to wake him. But even so, it is not hard to get the intuition that waking your dad might well be permissible. After all, the contrary claim — that it *must* be impermissible to wake him, because his well-being matters more than whether or not he approves — seems at least a little paternalistic. The fact that your dad will approve of being awoken cannot be dismissed merely because it runs counter to his own well-being. So it could make the difference to whether or not it is permissible for you to wake him up.

It will be helpful to describe this and other cases in terms of *reasons* and their *defeaters*, where a reason is “defeated” by a fact if, because that fact obtains, the reason is weaker than it would otherwise be. In the case at issue, you have a *pro tanto* reason to refrain from waking your dad, in virtue of the fact that waking him will make him worse off. But this reason is at least partially defeated by the fact that you are justified in believing that he will approve of having been awoken. On balance, then, it could be permissible for you to wake him. You will not exactly be doing him a *favor* by waking him — he needs his rest! — but you would not be doing anything wrong, either. The normative situation can be illustrated as follows:



Using a diagram here might seem a little self-indulgent, but it will be helpful as a starting point when the scenario gets more complicated. The thickness of the lines represents the strength of the reasons and defeaters, with lines of equal thickness representing reasons or defeaters of roughly equal strength. The diagram above illustrates that, in the absence of your defeater, you have a strong reason to refrain from waking your dad (because waking him will cause him considerable suffering), and a weaker reason to wake him (because you want to watch the meteor shower with him). Ordinarily, then, the balance of reasons would make it impermissible for you to wake him. But you are justified in believing that he will *approve*

of having been awoken, and this *defeats* your reason against waking him.⁹ That reason is weakened to the point that, on balance, it is permissible for you to wake him.

This is not to suggest that in any case with the same sort of structure as the meteor shower case, your action will be permissible. In fact, there are various ways of tweaking or filling in the details of the meteor shower story to construct a case in which you should not wake your dad. But I will argue that these “bad cases” are not like many cases of procreation. Many cases of procreation are like the “base case”: the case in which, intuitively, it *is* permissible for you to make your dad worse off by waking him up to watch the meteor shower. I will argue for this point by cataloging some “bad-making features” of meteor-shower-style cases: features in virtue of which the bad cases are bad. Then, having cataloged these bad-making features, I will argue that many cases of procreation do not share them.

3.1. *Lacking Justification*

The first bad-making feature is simple: it might be that, although your dad would approve of being awoken, you are not at all justified in believing this. Thus far I have assumed that, in the base case, your defeater for your reason against waking your dad is the fact that you are justified in believing that he will approve of having been awoken. So, in the bad case in which you lack justification, you lack a defeater for your reason against waking your dad. If you wake him, you are acting contrary to your undefeated reason against waking him, and your act is simply impermissible.

It is worth mentioning that my interpretation of the base case might strike some philosophers as overly subjectivist. Surely, the thought goes, what defeats your reason against waking your dad is not that you are *justified in believing that he will approve* of having been awoken, but simply that *he will approve*. This anti-subjectivist interpretation makes it a bit harder to explain what is wrong with waking your dad when you lack justification for thinking he will approve. After all, the anti-subjectivist cannot say that it is impermissible to wake your dad when you have a true-but-unjustified belief that he will approve of having been awoken, for in such a case your reason against waking him is defeated in the same way that it is in the base case. They can, however, say that waking him is irresponsible, in the sense that it reflects badly on you and you ought to be chastised for it. So they can vindicate the thought that there is something bad about your waking him in that case, even if what you do is technically permissible.


⁹ Notice I am not claiming that your dad’s future approval is a positive reason to wake him. It may be, but my central claim is that it is a *defeater* for your reason *against* waking him. Thus, my treatment of the case differs importantly from Elizabeth Harman’s treatment of similar cases. Harman writes “...the fact that if one performs an action then one will be glad one did it is a reason — though a defeasible reason — to believe that one should do it” (2009 p.194) From context, it is clear that Harman believes something similar is true of actions that affect others: *their* future approval is a defeasible reason to believe one should perform the relevant action. The difference between treating future approval as a *defeasible reason for action*, rather than a *defeater for a reason against acting*, is one which has implications for a number of the cases I will consider. While it would take me too far afield to track all of the implications, we can already see that there are implications for the meteor shower case. Suppose your dad will (strongly) approve of what you do, whether you wake him or not. If his future approval is a defeasible reason, as Harman argues, then his future approval at being awoken is a (strong) reason to wake him, his future approval at *not* being awoken is a (strong) reason to *not* wake him, and your (strong) well-being-based reason against waking him breaks the tie in favor of not waking him. So, on the Harman-style view, waking your dad is impermissible. If, on the other hand, his future approval is a defeater, then waking him is permissible. His future approval at not being awoken gives you no reason to not wake him, and his future approval at being awoken defeats your reason against waking him. So you are free to act on the fact that you want to watch the meteor shower with him. This is, I think, the more plausible treatment of the case.

I do not want to take a stand on the choice between these interpretations. Suffice it to say that if procreation turned out to be morally problematic in either of these ways — impermissible *or* irresponsible — then it would be bad news for procreators. So I will argue that procreation is not problematic in either respect.

3.2. Insufficient Defeat

A second bad-making feature concerns the balance of your reasons for and against waking your dad. It might be that, although your dad's future approval *partly* defeats the reason you have against waking him, what is left of that reason is still enough to seriously outweigh your reason for waking him.

To see this, suppose that waking your dad would cause him to suffer immensely. Whatever his ailment may be, waking him up in the middle of his sleep will make it much worse. His headache will become a week-long migraine, his nausea will prevent him from sleeping for days, and so on. You know that, even amidst this terrible suffering, your dad would approve of having been awoken. But if his suffering will be sufficiently excruciating, you should refrain from waking him nevertheless. The situation can be illustrated as follows:




The idea is that your reason against waking your dad is so strong that it greatly outweighs your reason to wake him, *despite* being weakened to a significant degree. Thus, you should overlook the fact that he will approve of being awoken, and refrain from waking him out of concern for his well-being. These claims have a distinctly paternalistic feel, but surely there is a point at which paternalism is appropriate. Your dad's future approval does not license you to inflict *unlimited* suffering. There is a point at which his suffering will be so terrible that you should not wake him even if he would approve.

It is not easy to say where the limit is exactly, but suffice it to say that the amount of suffering you can permissibly cause your dad is proportionate to the degree to which his approval is *serious* rather than *frivolous*. Seriousness and frivolity, in this context, are a matter of how much your dad genuinely cares about seeing the meteor shower. His approval might be *serious* if he is the sort of sensitive soul for whom scenes of natural beauty have profound or even quasi-religious significance. His approval might be *frivolous* if it is a mere whim that will be quickly forgotten. The former kind of approval, but not the latter, could make it permissible for you to make your dad significantly worse off by waking him. In other words, serious approval defeats your reason against waking your dad to a greater degree. In the base case, your dad's approval will be sufficiently serious (and the reduction to his well-being will be sufficiently modest) that his approval more or less fully defeats the reason you have against waking him, and so it is permissible for you to wake him.

3.3. Defeater Defeaters

A third and final bad-making feature concerns higher-order defeaters. Just as your dad’s approval can defeat the reason you have against waking him, his approval can *itself* be defeated by some further factor. To the extent that it is itself defeated, it does not reduce the weightiness of your original reason against waking your dad. In that case, the normative situation can be illustrated as follows:



Because your original defeater is itself defeated, the original defeater does little or nothing to blunt the force of your original reason against waking your dad. And because the force of that reason is preserved, it is impermissible for you to wake him. There are many possible defeater defeaters; I will consider two that are particularly relevant to the case of procreation.


First, your dad’s future approval might be defeated by your dad’s ignorance — it could be that if he knew more about the conditions of his awakening, then he would not approve of being awoken. Suppose he mistakenly believes — and will continue to believe — that losing sleep will not be detrimental to his health and will not leave him any worse off. If he knew more about the consequences of being awoken, he would not approve of being awoken. In this version of the case, too, your dad’s mistake is very plausibly a defeater for his approval. So the full force of your reason against waking him goes through, and on balance you should not wake him.

Second, your dad’s future approval might be defeated by its *inappropriateness*. Suppose your dad has a serious problem with self-loathing, and you are justified in believing that he will approve of being awoken precisely because he approves of his own suffering. This self-loathing-driven approval lacks the same kind of moral significance as the aesthetic-appreciation-driven approval from the base case. In the case in which your dad’s attitudes are masochistic, your actions should not be guided by respect for those attitudes, because they are wholly inappropriate. Here again, we have a defeater for your dad’s approval, so on balance you should refrain from waking him.

All told, we have three broad classes of bad-making features that might arise in meteor-shower-style cases. You could *lack justification* (you could fail to be justified in believing that your dad will approve of being awoken), your reasons could be *outweighed* (your dad’s approval could be insufficiently serious relative to the degree of his suffering) and there could be *defeater defeaters* (factors that make it the case that your dad’s approval is not a defeater for your reason against waking him). These bad-making features give rise to three corresponding worries for procreation. In the next section I will argue that each of the three worries does not apply to cases of procreation, or that it does apply but it can be overcome. So, in the end, many cases of procreation are like the base case. They are cases in which it is permissible to act in ways that make a person worse off, because that person will approve of your having done so.

4. Step Two: Potential Problems for Procreation

If the pessimistic hypothesis is true, then people are worse off for coming into existence. Thus, potential parents face a choice that is somewhat like the choice you face in deciding whether or not to wake your dad. They are choosing whether to act in a way that will make someone worse off, in light of their expectations about that person's future attitudes. Thus, a given case of procreation could share the same normative structure as the base meteor shower case:



The goal of this section is to show that many cases of procreation *do* have this structure, because they lack the bad-making features described in the previous section.

4.1. The Justification Worry

First consider the justification worry. In the meteor shower case, this was the worry that you might lack justification for believing that your dad will approve of being awoken. Similarly, procreators might lack justification for believing that their children will approve of having been born. If so, then they are doing something wrong by procreating.

Though there is little empirical data which deals *directly* with subjects' attitudes towards having been born, there is plenty of data that is at least suggestive. For example, there is a huge amount of psychological research concerning "subjective well-being", which is measured through such questions as "All things considered, how satisfied are you with your life as a whole nowadays?" (Diener 1987). In a compilation of seven large-scale surveys, covering 51 countries and 1.3 million randomly sampled people, most people reported being at least moderately satisfied with their lives (Blanchflower, & Oswald 2017). This in turn suggests that people approve of their own lives, broadly speaking.

Data regarding suicidality is also relevant to the question of whether people approve of their own lives. The data is sobering. For example: studies suggest that in recent years more than 15% of American adolescents had seriously considered suicide within the past twelve months (Xiao Y, Cerel J, Mann JJ 2021). Among adults the figure drops to around 4% (Ivey-Stephenson et. al 2022). It is at least plausible that many of those who suffer from suicidal ideation also sometimes wish that they had never been brought into existence. But this assumption is consistent with the claim that most people, most of the time, are satisfied with having come into existence. So it is consistent with the claim that many would-be procreators can be justified in believing that the people they create will approve of having been born.

None of this is to suggest that it is *easy* to be justified in the relevant belief. It might be more like being justified in believing that you will complete a difficult task, like an advanced hike. Not just anyone could be justified in having such a belief. One can only be justified if one prepares in advance, enjoys some favorable circumstances, and makes sincere commitments to see the task through. But if you have prepared,

your circumstances are favorable, you are sincerely committed, and the majority of hikers in similar circumstances have succeeded at similar tasks, then there is no reason you cannot be justified in thinking that you too will succeed. The same goes for procreation — there is no reason that many would-be procreators cannot be justified in thinking that the people they create will approve of having been created.

4.2. *The Outweighing Worry*

Just as your reasons to wake your dad might be outweighed by your reasons against waking him, one might worry that parents' reasons for creating new people are outweighed by their reasons against doing so. The idea is that, if the pessimistic hypothesis is true, we have a *very strong* reason not to procreate, and while future people's approval defeats that reason to a degree, on balance it is still strong enough that it wins out. So, the thought goes, procreation is a domain in which we ought to act paternalistically — ignoring future peoples' attitudes in the interest of preserving their well-being. To assess whether this worry can be overcome, we have to consider both the strength of our reason against procreating (in virtue of lowering peoples' well-being) and the strength of our defeater (in virtue of those peoples' future approval).

Consider first the strength of the defeater. Continuing with the assumption that the strength of the defeater is proportionate to the *seriousness* of the approval, then presumably the defeater in this case is very strong. Most of us care very deeply about our own lives. Insofar as we approve of having come into existence, this approval is not based on a mere whim, like the way in which you might approve of having a particular flavor of ice cream on some one-off occasion. On the contrary, concern for our own lives is among the deepest concerns most of us have. Presumably most of us care about our lives in a non-derivative way — we are simply glad to exist, in a way that admits no further rational explanation. But we also care about our lives in ways that are derived from our other serious concerns. Our having come into existence is a necessary precondition for much of what we care most deeply about — all of our accomplishments, projects, and relationships — and these serious concerns provide additional grounds for caring about having come into existence.¹⁰ So, if the fact that we approve of our own existence does not count as a “serious” instance of approval, then nothing does.

If, as in the meteor shower example, the amount of suffering that can be permissibly inflicted upon us is proportionate (all else being equal) to the seriousness with which we approve of the action that causes the suffering, then our parents were permitted to inflict a great deal of suffering on us by creating us. The same goes for future people. If they are anything like us, then their approval at having come into existence permits procreators to inflict, by creating them, a great deal of suffering on them.

With these observations at hand, we can see what it would take for the outweighing worry to be decisive against procreation. Suppose that some potential parents are considering having a child, Asha. Asha's life will matter a lot to her. Given the seriousness of Asha's approval, paternalism is permissible only if being brought into existence would cause an *extremely serious* reduction in Asha's well-being. It is hard to speculate about what sorts of lives might cause a sufficiently serious reduction in well-being. If Asha's parents can foresee that Asha's life will be pure unmitigated agony, that might seem to be enough to justify paternalism. But if Asha's parents can foresee *that*, then they should also have doubts about whether Asha

¹⁰ I do not mean to suggest that if one approves of one's friendships (for example) then one must approve of having come into existence. Someone could approve of having friends *given* that they came into existence, but disapprove of existence on the whole, including their friendships that go along with it. That would be a consistent — though I think somewhat unusual — set of attitudes. My point is that, if one *does* approve of one's existence on the whole, then one's approval of their friendships provides further grounds for approving of having come into existence.

will approve of having been brought into existence — in which case the real worry is the justification worry, not the outweighing worry. So it is highly unlikely that considerations of Asha's future well-being *alone* make it impermissible to create her.

4.3. The Defeater Defeater Worry

Finally we come to the worry regarding defeater defeaters. This is, I think, the most interesting and challenging of the three worries, although I believe it can be met.

First there is the case in which your dad is ignorant of the fact that waking him made him worse off. His ignorance defeats his approval insofar as he approves only because he lacks this relevant information. Similarly, future people could be ignorant of relevant information regarding their own existence — in particular, that being brought into existence was bad for them — and their ignorance could defeat their approval insofar as they approve only because they lack relevant information. The key claim here is *psychological*: people *would in fact* have different attitudes towards their own existence, if they had additional relevant information.

Second, there is the case in which your dad's approval at being awoken is somehow inappropriate. Insofar as his attitude is inappropriate — he really *should not* approve of being awoken — the inappropriateness defeats his approval. Similarly, if it turns out to be inappropriate for future people to approve of having come into existence, then the inappropriateness defeats their approval in the same way. The key claim here is *normative* — it is some sort of *mistake* to approve of having come into existence, assuming that the pessimistic hypothesis is true.

Putting these two worries together, the thought is that if the pessimistic hypothesis turned out to be true, and if we believed this, then we would react in one of two ways. Either we would stop approving of our existence, or we would stubbornly go on approving — but this latter reaction would be inappropriate. Either way, our attitudes of approval towards existence, as well as the approving attitudes of future people, are defeated. They lack the kind of normative authority I have attributed to them; they do little or nothing to blunt the force of our reasons against making future people worse off by bringing them into existence.

These worries can be made especially pressing with an argument from evolutionary psychology. According to this argument, we can reasonably expect that evolution would push us in the direction of thinking that our existence, and the continued existence of our species, is a good thing. There are various psychological data which appear to corroborate this reasonable expectation. David Benetar and Jason Marsh have argued that we have various optimistic biases which color our evaluations of our own lives (2006; 2014). Some psychological research suggests that we tend to recall good events more readily than bad events (Myers and Diener 1997), and we attribute meaning to negative events in a way that seems objectionably *post hoc* (Bering and Bjorklund 2004; Marsh 2014). These biases could explain why we might find it hard to accept the pessimistic hypothesis, and why we might go on approving of our own lives even if we did accept it. But, the thought goes, these are after all *biases*, and if our approval of our own lives is merely a product of our biases, then they lack the kind of normative authority which I have attributed to them.¹¹

These are important worries, but they can be met. The first thing to note is that biases pose more of a threat to the authority of cognitive attitudes like *belief*, as opposed to non-cognitive attitudes like *approval*. Suppose you're at a gymnastics competition in which your daughter is competing. You are, of course,

¹¹ I thank an anonymous reviewer for raising this worry.

biased towards your daughter, and this bias affects your attitudes in two ways: first, you believe that her routine was *better* than any other gymnast's routine; second, you *approve* of her routine more strongly than any other gymnast's routine. Your bias is more a threat to your belief than it is to your approval. Because you are biased, your belief is open to the objection that it is not truth tracking: you might believe that your daughter's routine was best when in fact it was not. But your bias does not expose your approval to the same kind of objection, because approval does not aim at the truth. Your bias might expose your approval to a different kind of objection: it might turn out that you approve of your daughter's routine *only because* you believe — as a result of your bias — that her routine was best. But if your approval is not contingent upon your biased belief in this way, then this objection doesn't stick either. Your approval is the result of bias, yes, but not *objectionably* so. It is not *inappropriate* for you to approve most strongly of your daughter's routine.

If this is right, and we can approve of things in a way that is unobjectionably biased, then in particular our attitudes towards our own lives might be unobjectionably biased. It might be that although our approval of our lives is the result of bias, this bias does not make our approval inappropriate. This is, I think, what we should say about our own case. We are biased, but our biased approval towards our lives is not inappropriate. The point can be illustrated with the following story:

Cradle Swap. You were quite nearly switched with another child at birth, in which case you would have been raised by different people and lived a very different life. That very different life would have been better for you than your actual life. Your hobbies would have been more fulfilling than your actual hobbies. Your achievements would have been more significant than your actual achievements. Your friends and family would have been at least a little cooler and kinder than your actual friends and family. So your other life would have been significantly better for you, though it would also be very different. Meanwhile, your actual friends and family would be no worse off (perhaps they would raise the child with whom you were swapped at birth, and the resemblance is such that they would never learn the difference.)

For any plausible theory of well-being, we can fill out the details of the Cradle Swapped life so that it is better for you than your actual life. If pleasure, achievement, meaningful relationships, or value fulfillment are good for us, then the Cradle Swapped life ranks higher than your actual life with respect to these goods. So we can be sure that, if you lived the Cradle Swapped life, you would be better off. Given that, the question is: do you wish that you had been swapped at birth, and lived the better life? Or do you approve of the fact that you were not swapped, even though you would be better off if you had been swapped?

Speaking for myself, I do not wish that I had been swapped at birth. I approve of the fact that I was not swapped at birth, but instead lived my actual life with my actual friends and family. These attitudes are unquestionably the result of bias: I am biased towards my actual friends and family, as well as my actual hobbies, accomplishments, etc. But it hardly seems objectionable to be biased in this way. It would be a different story if I approved of my life *only because* I (falsely) believed that my actual friends and family are better than the friends and family I would have had if I were Cradle Swapped. But I don't believe this. I know that I could have had some other friends and family who would have been better for me than my actual friends and family.¹² I know that the Cradle Swapped life would have been better than my actual life.

¹² With apologies to Patricia Pendoley and Mike Pallies, the best parents ever.

Even so, I do not wish that I had been swapped at birth, and I do not think that it is at all inappropriate to feel this way, despite the fact that my feelings are the result of bias. And I predict that this reaction will be widely shared. Most of us, when we think about cases like Cradle Swap, do not wish that we had been swapped.

In a recent, instructive paper, Camil Golub offers a similar diagnosis of what he calls our “conservative attitudes” towards our own lives (2019). The basic story is simple: we care about our actual relationships with friends and family, as well as our actual achievements, actual hobbies and interests, actual personal histories, etc.¹³ None of those features of our lives would be in place if we had been swapped at birth, so the other lives we might have lived would be missing much of what we actually care about.¹⁴ They would be missing much of what we actually desire. To be sure, swapped lives would satisfy many of our actual desires — we desire to have *some friendships or other*, for example, and those desires would be satisfied in our swapped lives. But we also desire to have the *actual friendships we have*, and those desires would not be satisfied.

There are a few points that need to be clarified here. First, I am not suggesting that our actual lives rank higher than our swapped lives with respect to desire satisfaction. It is true that we have satisfied *de re* desires for particular things (e.g. friendships and relationships) and we would not enjoy these *de re* desire satisfactions if we had been swapped at birth. But it’s not as if my actual life is the only one in which I would have satisfied *de re* desires for particular relationships. If I had been swapped at birth then I would have *different de re* desires for *different* friendships, and those desires would be satisfied. In fact, our swapped lives might contain a greater balance of desire satisfaction than our actual lives — that is just one of many respects in which our swapped lives might be better for us! But, crucially, our actual lives rank higher with respect to *actual-desire* satisfaction. With respect to our *actual* desires, the actual world is better, because it includes the actual friendships, achievements, etc., that we care about. This is why we approve of our actual lives over our swapped lives, and there is nothing inappropriate about our doing so.

The second point is that, although I am claiming that it is not inappropriate to approve of our actual lives over our Cradle Swapped lives, this is not a claim about the attitudes we might have had *before* we were swapped. If we were souls in Heaven, waiting to be incarnated on Earth, and we were given the choice between an ordinary life and a superior Cradle Swapped life, then presumably it *would* be inappropriate to choose the ordinary life. If we had not yet been born, and if we had no particular attachment to the things we’d do and the people we’d meet in either life, then we would have no grounds preferring the ordinary life over the Cradle Swapped life. But we do not occupy this lofty position — we are here, living on Earth, attached to our projects and relationships — so we *do* have grounds for approving of our ordinary lives over the Cradle Swapped lives. The fact that we *occupy* a particular possible world makes all the difference to which possibilities it is appropriate or inappropriate for us to prefer.

¹³ Golub offers two separate explanations for our conservative attitudes. First, we are attached to particular things in our actual lives. Second, we are attached to our “biographical identities” — roughly, the series of events in our personal histories that have shaped who we are as people (2019, p.79-82). I accept these explanations, but I prefer to roll them together and say that our “biographical identities” are simply another item on the list of particular things in our lives that we care about.

¹⁴ Elizabeth Harman makes a similar point, noting that events in our lives that make us worse off can lead us to form an *attachment* to features of our lives “...which makes a preference reasonable although it would otherwise be unreasonable” (2009 p.193).

One might worry that this last claim runs afoul of standard decision theory. Translating the claim into the language of preferences,¹⁵ rationality, and utility, the claim is that it is not irrational for subjects to prefer worlds that have lesser utility (even by the subject's own lights). The worry is that if utility is disconnected from rationality requirements in this way, then it — and by extension, standard decision theory — cannot do the philosophical work that it is supposed to do. And this seems like a major cost; we should not give up on the philosophical usefulness of decision theory unless we have very strong reason to do so.¹⁶

Luckily, the worry can be assuaged. We can retain the framework of decision theory while accepting the claim that it is not irrational for subjects to prefer worlds with lesser utility. We need only distinguish between closely related concepts that both go by “utility”: call them *objective utility* and *preference utility*. Objective utility is a matter of what is best; it is a function from subjects to worlds. Preference utility is a matter of what a subject rationally prefers, and is a function from *pairs of worlds and subjects* to worlds. This framework retains the descriptive power of decision theory, but it allows that the world one occupies can make a difference to what it is rational for one to prefer (but not to what is best). Given the fact that we are enmeshed in our lives, it is not irrational for us to prefer our lives over the very different swapped lives. We are biased towards our actual relationships (for example), this bias is not irrational, and as a result of this bias we prefer a life that does not lack these relationships.

These claims have straightforward implications for our attitudes towards having come into existence. If the pessimistic hypothesis is true, then we would be better off if we had not come into existence. But if it is not inappropriate to approve of the fact that one was not swapped at birth (despite the fact that having been swapped would be better for one) then it is not inappropriate to approve of having been born (despite the fact that not having been born would be better for one). In both cases, our approval is the product of our bias towards particular features of our own lives: our relationships, achievements, etc. Given that we are biased in these ways, and given that evolutionary forces have likely given us biases towards life itself, it is to be expected that we approve of our lives over non-existence. But if it is not inappropriate to be biased in one's approval of an ordinary life over a Cradle Swapped life, then why would it be inappropriate to be biased in one's approval of an ordinary life over non-existence? The considerations are the same in each case. The mere fact that our approval is biased does not mean that our approval is in any way objectionable.

To be sure, there could be creatures who would lack our biases, and who would form different judgments. And not even all humans will necessarily be biased toward approval. *Pure egoists*, with their clear-eyed and single-minded focus on well-being, would wish that they had been swapped at birth. *Pessimistic pure egoists* would wish that they had never been born at all. I am not suggesting that these

¹⁵ This worry only arises if one is willing to countenance the idea of preferences for states of affairs over which one has no control — like the fact that one was not swapped at birth. Even so, the worry can be re-introduced with a different and somewhat more contrived case. Suppose that there is a Twin Earth, and the cradle swap happens to your twin on Twin Earth. The question is whether you would press a button that causes you to swap places with your twin — swapping memories as well — so that you go on to live their superior life. (They are willing to make the swap for whatever reason, so there is no need to worry about the morality of resigning them to your inferior life.) I claim that it would not be irrational to refrain from pressing the button, even if you think that pressing it would be objectively better for you (and worse for no one).

¹⁶ Thanks to John Hawthorne for raising this worry and suggesting the line of response I sketch in the next paragraph. If the response does not work, it is my fault and not his.

attitudes are inappropriate, but our attitudes are not inappropriate either.¹⁷ We can approve of different things from pure egoists without either of us being wrong. Approval, unlike belief, leaves plenty of room for disagreement without error.¹⁸

This is not to say that we could *never* err in our approval of existing. It seems like we could in at least two ways. First, if our lives are *massively* worse than nonexistence, then presumably it is inappropriate to approve of them over nonexistence. But if, as seems plausible, our lives are less-than-massively worse than nonexistence (like the less-than-massive difference between our actual lives and Cradle Swapped lives) then this difference does not render our approval inappropriate.

Second, our approval could be mistaken if we only approve because we believe something false. In this case it could be that we approve of our lives only because we do not accept pessimism, and despite the fact that pessimism is true. This takes us at last to the worry about ignorance: the worry that we approve of our own lives only because we do not believe that the pessimistic hypothesis is true. So now we need to ask: assuming pessimism is true, would we go on approving of our own lives if we knew that pessimism is true? Or do we approve of our lives only because we are ignorant of relevant information about our lives — namely, they are worse for us than nonexistence?

If the foregoing arguments establish that it is not inappropriate to approve of what is worse for us, then this conclusion bears importantly on the question of whether we would approve of our lives if we had more relevant information about our lives. This is because, as a general rule, our attitudes track our beliefs about what sorts of attitudes are appropriate. For example, suppose you learn that your friend has recently saved five people from a serial killer, and naturally you approve of her having done so. She's a hero! Then you learn that, to save the five from the murderer, she had to kill one innocent person herself. You *may* continue to approve of what she has done, but that would seem to depend on whether you think it is appropriate to approve of killings that are performed as a means to prevent more killings. If you are enough of a consequentialist, you probably will continue to approve, at least on reflection. If you have a sufficiently deontological streak, you will probably cease to approve. Crucially, then, it's not *merely* learning about the killing that triggers a change in your attitudes of approval. Whether or not you cease to approve also depends on your background beliefs about the *appropriateness* of approval.

If our attitudes regarding existence follow the same pattern, then it seems people would continue to approve of their own lives. We are supposed to imagine that future people approve of having come into existence, but then learn that coming into existence was bad for them. Whether or not this would change their attitudes would seem to depend on their background beliefs about whether or not it is inappropriate to approve of their having come into existence (and more generally, of events that are bad for them). But if the foregoing arguments are sound, then it is *not* inappropriate, and we have to imagine that this is part of what future people would learn — for it, too, is relevant information. If changes in their attitudes of approval depend on their beliefs about the appropriateness of approval, then their attitudes would not change. They would react like consequentialists who learn that a very good state of affairs was brought about by a

¹⁷ One could argue that it would be inappropriate to wish that one was swapped, because one *should* approve of one's actual relationships, achievements, etc. I tend to disagree, but I remain officially neutral on this point.

¹⁸ As mentioned above, Jason Marsh argues that our beliefs about our own well-being are subject to optimistic biases, and takes this to pose a problem for procreation (2014). But he raises this challenge only tentatively, concluding "it is my hope that the challenges raised here can be answered in future work" (2014: f.n.39). I hope to show that the challenge can be met: optimistic biases threaten the normative authority of our beliefs, but not our attitudes of approval, and it is these latter attitudes that are more relevant to the ethics of procreation.

somewhat bad state of affairs: they would continue to approve because they believe that the new fact is irrelevant to whether or not approval is warranted.

Ultimately, it is an empirical question whether gaining additional relevant information would cause future people to stop approving of having come into existence. But given that it would not be inappropriate for them to approve, and given that this is itself relevant information, it is hard to see why their gaining relevant information would lead them to stop approving. So we should provisionally conclude that not only is it not inappropriate to approve of having come into existence, but also that people would approve of having come into existence — even if they accepted pessimism. Putting these conclusions together, we should conclude that future peoples' attitudes are not defeated by either inappropriateness or ignorance.

5. Paternalism and the Benefits of Existence

I am not the first to suggest that the attitudes of future people could be relevant to whether or not it is permissible to create them. Drawing upon an earlier paper by Seana Shiffrin (1999), Asheel Singh has recently argued that would-be procreators cannot appeal to the hypothetical consent of future people to justify procreation (2018). But the views that Singh criticizes are importantly different from my own view, in ways that turn out to be instructive.

Singh's official target is the view that procreation is permissible in virtue of the hypothetical consent of future people, but he takes there to be a strong connection between this view and a certain kind of *paternalism*. He writes:

I shall use *paternalist* and *paternalism* when referring to those who believe that certain harmful acts can be justified via appeals to hypothetical consent—such is the strength of the link I see between a defense of paternalism and an appeal to hypothetical consent. (2018, p.1140)

At first blush, it seems odd to see such a strong link here. The question is whether we should care more about future people's *well-being*, or future people's *attitudes*. It seems *anti-paternalistic* to care more about their attitudes, rather than their well-being. That is why I have previously suggested that it is paternalistic — though not necessarily wrong — to refrain from making someone worse off despite knowing that they would approve of your having done so.

The link Singh sees here becomes clearer in the context of his view about what hypothetical consent amounts to. He writes that hypothetical consent comes in two varieties: *subjective* hypothetical consent, and *objective* hypothetical consent. The former is assessed by asking: "To what sorts of proposals would it be rational for *agent A* to consent, given *her* unique aims?" (2018, p.1142). The latter is assessed by asking: "To what sorts of proposals would contractors in the 'original position' (Rawls 1971) be most reasonable to consent?" (2018, p.1141). Singh contends that we cannot rely on subjective hypothetical consent in the case of future people, roughly because future people do not presently have any unique aims and their would-be parents cannot predict what their future aims will be (2018, p.1142). This only leaves objective hypothetical consent, which would-be procreators can secure only by providing certain *objective benefits* to people they create, such that the total package of benefits-plus-harms would be chosen by contractors in the original position. So, for Singh, appeals to hypothetical consent go hand in hand with attempts to justify harming future people by claiming that those people objectively ought to accept the package of benefits-plus-harms that existence brings. That *does* seem paternalistic. Returning to the meteor shower case, it is tantamount to saying that it is permissible to wake your dad because he *should* approve of being awoken,


because he should recognize that the benefit of seeing the meteor shower outweighs the hit to his well-being.

I agree that this brand of paternalism is not a promising strategy for justifying procreation. But this has not been my strategy. On my view, what matters in the first instance is not whether future people *should* approve of being created, but whether we are justified in thinking that they *will*. For reasons I canvassed in the previous section, I am not nearly as skeptical as Singh regarding our ability to predict the attitudes of future people. We do not need to know their “unique aims”; we only need to be justified in believing that they will approve of having come into existence. The question of whether or not future people will be benefitted by their creation plays no role in my argument; indeed I have been assuming that on the whole people are worse off for coming into existence.

Singh could point out, however, that benefits do play a kind of indirect role in my argument. I have said that if future people’s approval of having come into existence is *inappropriate*, then the inappropriateness of their attitudes is a defeater for those attitudes. And presumably, if the benefits one receives by coming into existence are sufficiently meager in comparison to the reduction of one’s well-being, then it is inappropriate for one to approve of having come into existence. If one’s life is unmitigatedly miserable, then presumably it would be inappropriate for one to approve of having been brought into existence. Such an existence would include nothing, or very little, that is worth wanting. So, Singh could argue, I am at the end of the day committed to the view that procreation is justified only if the harms of existence are sufficiently compensated with benefits. The conclusion would then be that my justification of procreation is paternalistic in precisely the same way as Singh’s targets, albeit with some camouflage over the paternalism.

I agree that, on my view, procreation is justified only if the harms of existence are in a certain sense compensated with benefits. But I deny that this makes my view paternalistic in the same way as Singh’s targets. There are two related reasons for this. First, the point at which the harms are sufficiently “compensated” by benefits is not the point at which the created person *ought* to accept the total package of benefits-plus-harms. Whether we understand benefits purely in terms of well-being, or in terms of other kinds of value, I do not claim that procreation is justified only if the benefits of existence *outweigh* the reduction to one’s well-being that comes with being brought into existence. More generally, the idea that future people ought to approve of their having come into existence plays no role in my arguments. For all I have said, it might be that it is *never* inappropriate to disapprove of having come into existence.

Second, and relatedly, while the view I have described gives a role to future benefits, that role is not to *justify* procreation. The fact that a given future person will enjoy certain benefits — they will have friends and family, for example — means at most that we *lack* a certain reason for doubting the significance of their future approval. It means the normative situation is not one in which there is a defeater defeater:



Future benefits have the role of “crossing out” the defeater defeater box in the above diagram, so the original defeater goes through. Future benefits do *not* have the role of supplying a positive reason in favor of creating the child. So there should be no concern that procreation is being defended on paternalistic grounds, by appealing to benefits of which future people “should” approve.

6. The Procreation Asymmetry and the Procreation Asymmetry

I have argued that, contrary to popular belief, the pessimistic hypothesis does not have dramatic implications for procreation. Many future people will approve of having come into existence, and for the most part their attitudes towards their own existences will not be inappropriate, despite the fact that they are worse off for having come into existence. Many would-be procreators can be justified in believing that the people they create will approve of having come into existence, and as a result they can be justified in procreating even though the people they create will be worse off for having been created.

In this way, I have tried to show that the pessimistic hypothesis does not clash with our intuitions regarding procreation. In closing, I want to briefly note that the pessimistic hypothesis actually does better than optimism with respect to some other intuitions regarding procreation. In particular, the truth of the pessimistic hypothesis could explain our intuitions regarding the so-called “procreation asymmetry.” It’s often said that while we are obligated to avoid creating people who will lead bad lives, we are *not* obligated to create people who will lead good lives (Naveson 1967; McMahan 1981; Earl 2017). The challenge is to explain why our obligations are asymmetrical in this way.

The conclusions of this paper suggest a way in which this pattern in our obligations might be explained. If the pessimistic hypothesis is true, then no one will lead a good life, in the sense of a life that is good *for them*. So it is easy to explain why we are not obligated to create people who will lead good lives — there are no such people. And the foregoing arguments of this paper, if sound, can explain why it is permissible to create some people but not others. The ones we are permitted to create are those who we can be justified in believing will approve of having come into existence. The ones we are not permitted to create are those who we cannot be justified in believing will approve of having come into existence. Those in the latter camp are primarily those people whom we can be justified in predicting will grow up in unmitigatedly painful, traumatic, or terrible conditions. Thus, we are sometimes permitted but never obligated to create new people.

To be clear, this explanation would not vindicate the procreation asymmetry as it is ordinarily formulated. Indeed, it is inconsistent with the standard formulation, since it entails that it is sometimes permissible to create people whose lives are bad for them. But it does vindicate the intuition that it is

permissible (but not obligatory) to create people who will lead *basically normal* lives, and impermissible to create people who will lead *terrible* lives. This is, presumably, the intuition that motivates standard formulations of the procreation asymmetry. And if pessimism is true, then the standard formulation of the procreation asymmetry is simply a failed attempt to articulate our intuitions about procreation. The mistake is in assuming that the distinction between normal lives and terrible lives is the same as the distinction between lives that are good for those who live them and lives that are bad for those who live them. Once we drop this assumption, and embrace the view that all lives are to some degree bad for those who live them, our intuitions become much easier to explain.

I conclude that procreation — or at least, many cases of procreation — are permissible, even if the pessimistic hypothesis is true. Many would-be procreators can be justified in believing that the people they create will approve of having been created, and their approval is sufficient to defeat the reason that the would-be procreators have against creating someone who will be worse off for having been created. Furthermore, there is no compelling reason to doubt the moral significance of future peoples' attitudes. In particular, there is no compelling reason to think that their attitudes will be either inappropriate or problematically ill-informed. In many cases, then, it is on balance permissible to create new people — even if their lives are bad for them.

Works Cited

- Benatar, David (2006). *Better Never to Have Been: The Harm of Coming Into Existence*. New York ;Oxford University Press. <https://doi.org/10.1093/9780199296422.001.0001>
- Bering, J. M. and Bjorklund, D. F. (2004). The Natural Emergence of Reasoning about the Afterlife as a Developmental Regularity. *Developmental Psychology* 40(2), 217–233. <https://doi.org/>
- Biswas T, Scott JG, Munir K, et al. Global variation in the prevalence of suicidal ideation, anxiety and their correlates among adolescents: A population based study of 82 countries. *EClinicalMedicine*. 2020;24:100395-100395. doi:10.1016/j.eclinm.2020.100395
- Bramble, Ben (2013). The distinctive feeling theory of pleasure. *Philosophical Studies* 162 (2):201-217. <https://doi.org/10.1007/s11098-011-9755-9>
- Campbell, Stephen M. & Nyholm, Sven (2015). Anti-Meaning and Why It Matters. *Journal of the American Philosophical Association* 1 (4): 694-711. <https://doi.org/10.1017/apa.2015.9>
- Chalmers, David (2002). The components of content. In David J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press.
- Earl, Jake (2017). A portable defense of the Procreation Asymmetry. *Canadian Journal of Philosophy* 47 (2-3):178-199. <https://doi.org/10.1080/00455091.2016.1274867>
- Golub, Camil (2019). Personal Value, Biographical Identity, and Retrospective Attitudes. *Australasian Journal of Philosophy* 97 (1):72-85. <https://doi.org/10.1080/00048402.2018.1431297>
- Harman, Elizabeth (2009). 'I'll Be Glad I Did It' Reasoning and the Significance of Future Desires. *Philosophical Perspectives* 23 (1):177-199. <https://doi.org/10.1111/j.1520-8583.2009.00166.x>
- Heathwood, Chris (2019). Which Desires Are Relevant to Well-Being? *Noûs* 53 (3):664-688. <https://doi.org/10.1111/nous.12232>
- Hurka, Thomas (2020). The Parallel Goods of Knowledge and Achievement. *Erkenntnis* 85 (3):589-608. <https://doi.org/10.1007/s10670-020-00245-0>
- Ivey-Stephenson AZ, Crosby AE, Hoenig JM, Gyawali S, Park-Lee E, Hedden SL. Suicidal Thoughts and Behaviors Among Adults Aged ≥18 Years — United States, 2015–2019. *MMWR Surveill Summ* 2022;71(No. SS-1):1–19. <https://doi.org/10.15585/MMWR.SS7101A1>
- Kagan, Shelly (2014). An Introduction to Ill-Being. *Oxford Studies in Normative Ethics* 4:261-88. <https://doi.org/10.1093/acprof:oso/9780198722144.003.0012>

- Lawrence, H.R., Burke, T.A., Sheehan, A.E. et al. Prevalence and correlates of suicidal ideation and suicide attempts in preadolescent children: A US population-based study. *Transl Psychiatry* 11, 489 (2021). <https://doi.org/10.1038/s41398-021-01593-3>
- Marsh, Jason (2014). Quality of Life Assessments, Cognitive Reliability, and Procreative Responsibility. *Philosophy and Phenomenological Research* 89 (2):436-466. <https://doi.org/10.1111/phpr.12114>
- McMahan, Jefferson (1981). Problems of Population Theory: Obligations to Future Generations. R. I. Sikora, Brian Barry. *Ethics* 92 (1):96-. <https://doi.org/10.1086/292301>
- Myers, D. and Diener, E. (1997). Summary of Who is Happy? In Ackerman F., Kiron, D., Goodwin, N., Harris, J. and Gallagher, K. (eds.), *Human Well Being and Economic Goals*. Washington, DC: Island Press, 174–177. <https://doi.org/10.1111/j.1467-9280.1995.tb00298.x>
- Narveson, Jan (1967). Utilitarianism and new generations. *Mind* 76 (301):62-72.
- Rice, Christopher M. (2019). Objective List Theories and Ill-Being. *Ethical Theory and Moral Practice* 22 (5):1073-1085. <https://doi.org/10.1007/s10677-019-10035-5>
- Schwitzgebel, Eric (2006). The unreliability of naive introspection. *Philosophical Review* 117 (2):245-273. <https://doi.org/10.1215/00318108-2007-037>
- Shiffrin, Seana (1999). Wrongful Life, Procreative Responsibility, and the Significance of Harm. *Legal Theory* 5 (2):117-148. <https://doi.org/10.1017/S1352325299052015>
- Singh, Asheel (2018). The Hypothetical Consent Objection to Anti-Natalism. *Ethical Theory and Moral Practice* 21 (5):1135-1150. <https://doi.org/10.1007/s10677-018-9952-0>
- Sumner, Wayne (2020). The Worst Things in Life. *Grazer Philosophische Studien* 97 (3):419-432. <https://doi.org/10.1163/18756735-000108>
- Veenhoven, Ruut. (2020). "World Database of Happiness: a 'findings archive'". In *Handbook on Wellbeing, Happiness and the Environment*. Cheltenham, UK: Edward Elgar Publishing.