

Why We Should Create Artificial Offspring: Meaning and the Collective Afterlife

John Danaher¹

Received: 21 December 2016 / Accepted: 6 June 2017 / Published online: 3 July 2017
© Springer Science+Business Media B.V. 2017

Abstract This article argues that the creation of artificial offspring could make our lives more meaningful (i.e. satisfy more meaning-relevant conditions of value). By ‘artificial offspring’ I mean beings that we construct, with a mix of human and non-human-like qualities. Robotic artificial intelligences are paradigmatic examples of the form. There are two reasons for thinking that the creation of such beings could make our lives more meaningful and valuable. The first is that the existence of a collective afterlife—i.e. a set of human-like lives that continue after we die—is likely to be an important source and sustainer of meaning in our present lives (Scheffler in *Death and the afterlife*, OUP, Oxford, 2013). The second is that the creation of artificial offspring provides a plausible and potentially better pathway to a collective afterlife than the traditional biological pathway (i.e. there are reasons to favour this pathway and there are no good defeaters to trying it out). Both of these arguments are defended from a variety of objections and misunderstandings.

Keywords Meaning · Artificial offspring · Collective afterlife · Robotics · AI · Human enhancement

Introduction

This article defends an unusual thesis. An illustration might help. In season 3 of *Star Trek: The Next Generation*, in an episode entitled “The Offspring”, Lt. Commander Data (an artificial humanoid life-form) decides to create an artificial child of his own. He calls the child “Lal” and gives it the characteristics of a female human. The episode then tells the charming story of Data’s emerging bond with Lal, a bond

✉ John Danaher
johndanaher1984@gmail.com; john.danaher@nuigalway.ie

¹ School of Law, NUI Galway, University Road, Galway, Ireland

which is threatened by a group of scientists who try to take her away and study her, and eventually leads to heartbreak as Lal “dies” due to a design fault (though not before Data downloads her memories into his own brain so that a part of her can live on). The story is typical of the series. Throughout the seven seasons of its original run, the *Next Generation* writers repeatedly portrayed Data as a being that is trying to overcome his artificiality and become more human. “The Offspring” is a near-perfect example of this story-telling motif. What could be more human than the desire to have a child? The writers get us to sympathise with Data’s character in the episode by bringing to mind emotions and desires that are mirrored in human parent–child relationships.

The thesis here is that our sympathies in this case go in the wrong direction. We should not sympathise with Data’s attempts to become like us. Instead, we should hope to become more like him. We should hope to create artificial (non-human, but human-like) offspring of our own because doing so would be to commit to a project that could enhance the meaningfulness of our lives.¹ This thesis is defended with a two-step argument. Step one claims that the existence of a collective afterlife (Scheffler 2013) is an important source and sustainer of meaning in our lives. Step two claims that the creation of artificial offspring is both (a) a credible pathway to a collective afterlife and (b) a potentially superior pathway to a collective afterlife than our pre-existing biological pathway. Most of the attention is dedicated to defending the second step. But this is done after first explaining what is meant by the term ‘artificial offspring’ and outlining the logical structure of the argument being defended.

What are Artificial Offspring?

It is important to clarify what is meant by “artificial offspring”. It is a concept that may raise some hackles, but clarification is not an easy task. For starters, the concept of the “artificial” is contested. It is sometimes clarified by contrasting it with the *natural*,² and it is proposed to adopt this contrast, but it is admitted that the contrast

¹ Throughout this article, the view that meaningfulness is not an either/or property is adopted. The meaningfulness of one’s life is determined by certain conditions of value that one’s life does, or does not satisfy. The more such conditions that are satisfied, the more meaningful one’s life is. Consequently meaningfulness can vary both across and within lives. It is accepted, however, that a certain threshold of such conditions may need to be satisfied in order to make one’s life *worthwhile*, but that is taken to be a distinct concept from meaningfulness. This take on meaningfulness as a variable property is adopted by others in the literature (e.g. Metz 2013). It is also accepted that meaningfulness has something to do with the flourishing and survival of creatures like us, but that it is not limited in its application to the members of *Homo Sapiens* (i.e. that non-human creatures can have meaningful lives). This is a view that is defended, gradually, over the course of article, particularly in “[The Superiority of Artificial Offspring](#)” section where it is argued that phenotypic and not genotypic properties are what count when it comes to meaning.

² The term ‘natural’ is open to broad and narrow interpretations. A broad interpretation would hold that the natural is everything that exists in the natural (non-supernatural) world. This would encompass artificial, manufactured entities. A narrower interpretation would hold that the natural is anything that exists apart from human creation and manufacturing. This would distinguish the natural from the artificial. It is this latter, narrower, interpretation of the term that is intended in the text.

is not always helpful since the “natural” is also contested.³ On top of this, the idea of referring to artificial beings as “offspring” will be anathema to some people. If one uses any of these terms without supplying some interpretive guidance, one risks being misunderstood. The purpose of this section is to supply that guidance.

At a first pass, the “artificial” can be defined as that which is made by the application of a skill or technique (directly or indirectly) by a human being,⁴ i.e. that which is created from raw materials that are actively sourced and manipulated by human beings into some pre-planned shape or form (with some general goals or intended uses). This pre-planning need not anticipate every peculiarity or consequence that will befall the created artifact, but it must have some reasonably definite initial content. Thus, for example, the cup on my table is artificial because it was manipulated by human creators into a vessel for carrying other substances, even if those creators could not have anticipated that I would use it for holding my pens and pencils.

Second, it is submitted that there are paradigmatic and non-paradigmatic instances of artificial objects. The smartphone on my desk is a paradigmatic instance of an artificial object. The cows in the field outside my window are not. They lie on the contested border between the natural and the artificial. They are natural to the extent that they are (in large part) products of biological evolution. They are artificial to the extent that, their biological evolution notwithstanding, they have been selectively bred by humans, over thousands of years, to exhibit a certain set of characteristics. There are interesting debates one could have about the classification of such non-paradigmatic objects, some of which could take on metaphysical or ethical significance.

Third, and following on from this, it is submitted that there are paradigmatic and non-paradigmatic examples of artificial offspring. Data’s daughter Lal would (if she were real) be a near paradigmatic example of artificial offspring. She would be made from non-biological materials. She would be manufactured into existence by another artificially created being (himself created by a human) in a way that bypassed the ordinary mechanisms of biological conception and development. She would be contrasted with a paradigmatic example of a natural child, one that is conceived and developed through the ordinary biological mechanisms of conception and development. There are many interesting cases that lie on the border. Are children conceived in vitro (as opposed to in utero) natural or artificial? How about children that are genetically or biologically engineered? Again these are important questions and how we answer may have some ethical and axiological significance, some of which are discussed later in this article. In the meantime, it is upon the paradigmatic cases of artificial offspring that the argument is focused.

So much for the ‘artificial’ element of the concept. What about the ‘offspring’ element? Care is needed here. Some people will insist that the term “offspring” can only be applied to biological human children. The purpose of this article is to argue

³ This will be familiar to anyone who has debated the merits of natural law ethics, or who has entered into the nature-nurture debate.

⁴ Our own artificial creations could make artificial creations of their own. This would not move them out of the category of the artificial, even though the link to human agency is indirect. This is why Data’s creation of Lal falls within the scope of my definition of ‘artificial offspring’.

that this is too narrow an understanding of “offspring”. The argument is going to be that artificial beings, with the right mix of properties, can be functionally equivalent to human beings in terms of their relations to us as future generations. To anticipate that argument, and to help the reader understand the terminology being adopted, the claim is that in order for an entity to count as our offspring it is not essential that the entity be akin to a ‘child’ (i.e. have some special bond or relationship of dependency with a parent or need to go through some development and maturation) but that the entity in question is *human-like*, but non-human in nature. Understanding the right mix of human-like and non-human-like attributes is central to the argument that will be made. It is only artificial offspring with the right mix of attributes that would contribute to the meaningfulness of our lives.

What are these attributes? The precise specifications will emerge later on. Nevertheless, some general guidelines can be given. In terms of human-likeness, the artificial offspring should perceive and interact with the world in a similar manner to human beings. Furthermore, they should share a similar framework of values (how many is something that is open to dispute). That is to say: they should care about similar issues, have similar interests and be willing to pursue similar (but not equivalent) plans. A willingness to participate in individualised relationships with their human creators, analogous to those that subsist between parents and children might be an advantage, but perhaps not a prerequisite. In terms of non-human likeness, they should be free from many of the biological frailties that afflict human beings (for reasons I shall get into later). In particular, they should be well-protected from human-like forms of infirmity and disease. This does not mean that they have to be immortal and invulnerable. Far from it. It just means that they should be more existentially robust than ordinary biological offspring. Again, the human-like robots and androids of science fiction (such as Data)⁵ are probably paradigmatic in terms of their human-likeness and robustness.

This, then, is the thesis that will be defended in the remainder of this: creating artificial offspring, who are human-like in terms of their cognitive and evaluative frameworks but less prone to the biological frailties that afflict ordinary human beings could be a great source of meaningfulness in our lives. Indeed, creating them might be a greater source of meaningfulness than having ordinary human children.

How could one defend such a claim?

The Basic Argument

This is the argument that will be defended:

1. The collective afterlife (that is: the lives of all those who continue to exist on earth or elsewhere in the natural universe after we die) is an important source and sustainer of meaning in our lives. (*The Collective Afterlife Thesis*).
2. We have reason to pursue projects that have a greater probability of supporting or maintaining that which provides our lives with value and meaning.

⁵ Or, at least, Data with his emotion chip installed since emotion may be essential to sharing a similar evaluative framework to ordinary human being.

3. Therefore, we have reason to pursue projects that have a greater probability of ensuring that there is a collective afterlife.
4. The creation of artificial offspring has a greater probability of ensuring that there is a collective afterlife that sustains the conditions needed for meaning and value than the creation of traditional biological offspring.
5. Therefore, we have reason to pursue the creation of artificial offspring.

The first premise of this argument is taken from the work of Samuel Scheffler (2013) and will be explained in more detail below. The second is intended to state a prudential axiological reason for action. This is important. The argument should not be interpreted as claiming that we have some all-things-considered obligation to create artificial offspring. It is merely claiming that, to the extent that we care about living meaningful lives (or, rather, lives of *greater* meaning), we should seriously consider the project of creating artificial offspring. The possibility that we have competing moral obligations that outweigh the axiological wisdom of that claim will not be ignored. Indeed, it is something that is discussed in later sections when dealing with criticisms of the argument.

The crux of the argument lies in premise (4). This premise really breaks down into two sub-claims. First, it claims that the creation of artificial offspring is a credible pathway to the collective afterlife that Scheffler alludes to in his work. Second, it claims that it may be a superior pathway than the existing biological pathway. The combination of both claims is important. The first claim, in itself, states something pretty significant, viz. that the creation of artificial offspring could be a source and sustainer of meaning in our lives. But when combined with the second claim it becomes even more significant. It suggests that not only is the creation of artificial offspring *a* source and sustainer of meaning, it may actually be a *superior* source and sustainer of meaning.

If this article can succeed in defending premise (4), then the conclusion would seem to follow and it will have provided a serious reason to consider the project of creating artificial offspring. It is important to realise that the argument is not that we have strong reason to favour the creation of artificial offspring because that would make lives more valuable and more meaningful *in the future*. Rather, the claim is that we have strong reason to do it because it will make *our present lives* more meaningful. This is in line with Scheffler's claims about the collective afterlife, and it is important that it be emphasised. While it does matter what happens to the future generations of offspring (artificial or otherwise), the argument is not that it matters because of how it affects them. It matters because of how it affects us.

Now let's consider the premises in more detail.

A Modified Collective Afterlife Thesis

We must start with premise (1): the collective afterlife thesis. The goal here is not to offer a full-blown defence of this thesis. Scheffler has already done that in his work. Nevertheless, it is important to clarify what it means, to summarise some of the lines of argument that might lead one to endorse it, and to modify it slightly so that it fits the argument defended in this article.

A few definitions will be helpful, starting with the concept of a collective afterlife:

Collective Afterlife: The continued existence of human-like beings on this earth or elsewhere in the natural universe, in conditions roughly equivalent to those in which you now live, after your death.

This definition is mine, not Scheffler's. Three features of it should be noted. First, note how the collective afterlife being referred to has no spiritual or religious connotations. It is not an afterlife consisting in the continued existence of beings in a supernatural or heavenly realm. It is an afterlife involving the continued existence of beings on this earth, or elsewhere in the universe. It is consequently a wholly naturalistic, non-religious concept of the collective afterlife that is being appealed to throughout Scheffler's book and this article. Second, note how the definition refers to "human-like beings", not humans. This is one place where my definition differs from Scheffler's. In his book, Scheffler talks (or implies) that his collective afterlife involves the existence of human beings. The necessity of this is being disputed in this article. The argument here is that it is not absolutely essential for the beings that exist in the collective afterlife to be human (i.e. genetic members of *homo sapiens*). Human-like beings, with similar properties of personhood and similar goals and aspirations would be sufficient. This will be defended later on but it is important that the definition allow for it at the outset. Third, the final part of the definition is more purely Scheffler's view and claims that these future generations must live in conditions roughly equivalent to those in which we now live. The precise conditions in which future generations must live is somewhat contentious as between Scheffler and his critics (Johnston 2014; Kauppinen 2014). Nevertheless, it is pretty clear that, for Scheffler's argument to work, the lives of future generations must share at least some of our values, aspirations and needs.

Scheffler argues that the collective afterlife (understood in the non-religious, wholly naturalistic sense) is an important source of meaning and value in our present lives. He calls this the *dependency thesis*. As he puts it himself, our understanding of what makes a human life good (well-lived and meaningful) seems to be dependent on the assumption that it forms part of an ongoing pattern:

...our conception of "a human life as a whole" relies on an implicit understanding of such a life as itself occupying a place in a ongoing chain of lives and generations. (Scheffler 2013, 43)

If that chain of lives and generations came to an end, then our lives would be denuded of much of their meaning and value. This claim is simply the obverse of premise (1) of the argument being defended in this article: that the collective afterlife is a sustainer of meaning in this life. There are several ways in which it could do this. Scheffler, in particular, suggests that it could be because of *the actual existence* of a collective afterlife, or because of *the belief in* the general concept or idea collective afterlife (2013, 51–53).⁶ This article will presume an *actual existence*

⁶ To be clear: Scheffler initially defends a purely attitudinal version of the dependency thesis, but he later (2013, 51–52) suggests that the actual existence of a collective afterlife must be important to our conception of value. This is for the simple reason that taking a pill that induced false belief in the

version of the thesis, but note that the argument will work with a *belief* oriented version as well. This is because the argument is primarily about the *reasonable confidence* we can have in the future existence of artificial offspring, not in their existence per se. Nevertheless, one assumes that the reasonableness of our confidence must, in some way, be linked to the probability of their actual existence. This is why an actual existence version of the dependency thesis is favoured throughout the article, and why all references to the collective afterlife throughout the paper are not modified by some qualifier such as “the concept of” or “the belief in”.

This is enough definitional clarification. What of the actual defence of the collective afterlife thesis? Scheffler uses two thought experiments to support it. They will be briefly outlined here and the intuitive reactions and philosophical insights that Scheffler presumes we should draw from them will be highlighted. The intuitive reactions turn out to be key. If you do not share the intuitive reactions, you are unlikely to be moved by the remainder of the argument in this paper. But those intuitive reactions will not be defended in depth here. I feel them, and I will presume for the sake of argument that the reader does too. This is what helps to motivate the remainder of the argument. If you do not share them, you may still find it interesting to see how it is possible to link the thesis about artificial offspring to Scheffler’s thesis about the collective afterlife, but you are unlikely to be moved by the argument at a deeper, reason-giving, level.

The first thought experiment is:

The Doomsday Thought Experiment: Suppose that you will live a long, normal human life, but that 30 days after your death, all human life will be destroyed in some catastrophic event (for example, an asteroid collision). Suppose, further, that you know this catastrophic event will take place as you are living your life. What effect would this have?

Scheffler suggests, in a lengthy analysis (2013, 18–37), that it would have a pretty devastating effect on your life. It would rob many of your projects and activities of their value, and would probably induce a significant amount of despair, grief and existential hand-wringing. He suggests that our reactions to such a scenario tells us two important things about what makes our lives worth living. The first is that there appears to be a strongly *nonexperiential* aspect to what makes life worth living.⁷ In the doomsday scenario, your life and experiences are unaffected—you do not die prematurely—but nevertheless the value of your life (or at least your attitude toward it) is, somehow, affected. The second thing is that our reaction suggests that there is a significant *conservatism* to what makes our lives valuable. In other words, part of the reason why we care about certain things now is that we want them continue to be relevant after we die. I could spend my life searching for a cure for cancer (and

Footnote 6 continued

existence of a collective afterlife would not (prospectively) convince us that our lives retained the value we previously thought they possessed.

⁷ Similar to the case made by Nozick in the Experience Machine thought experiment. Nozick, R. *Anarchy, State and Utopia* (Harvard University Press 1974).

succeeding), or producing some great work of art. This could imbue my life with great meaning. But some of that meaning would be robbed if it turned out that, after I was gone, no one would have need for my cure or would be around to appreciate my art. Combined, these two implications provide support for the *dependency thesis*. We need to have (or believe in) the existence of future generations of beings that are *sufficiently like us* in order to sustain much of the meaning and value of the lives we presently live.

One problem with the doomsday thought experiment, however, is that it may conflate the importance of the continued existence of beings with lives that are *close* to our own with the continued existence of beings with lives *like* our own. It could be, for all the doomsday thought experiment suggests, that what induces the despair and existential angst is the fact that our children, friends and family, or any other being close to us, will die. Although Scheffler thinks the continued existence of such beings is an important part of what confers value on our lives, he thinks that their existence alone does not do justice to the dependency thesis. This leads to the second thought experiment:

The Collective Infertility Thought Experiment: Suppose that the entire human race is infertile. In other words, the current generation of humans is the last generation of humans that will ever live. What effect would that have on our lives?

Again, Scheffler suggests that it would have a pretty devastating effect (2013, 37–51). It would induce a significant amount of despair and existential angst. This is something that the book and film *The Children of Men* tries to illustrate in some rich, imaginative detail. The despair in the collective infertility scenario is not just caused by the prospective deaths of ourselves and people we care about. The despair in the collective infertility scenario is caused by the fact that *everyone*—including those with whom we have no special or personal connection—is gradually going extinct. The fact that we feel despair at this generalised extinction tells us something interesting. It tells us that there is a strong *altruistic* element to the role of the collective afterlife in our own lives. We care about the *general fate of beings living lives like ours*, not just the fate of people we know and love.

This adds further support for the dependency thesis and is essential to the argument in the present article. One thing that the infertility thought experiment seems to suggest is that biological/genetic parenthood over the future generations is not the only thing that helps those generations sustain the meaningfulness of our present lives. Rather it is some set of general characteristics, shared by the entire population (past and future), that sustains the meaning and value. What is this set of characteristics? Some might argue that the relevant general characteristics are that we are all members of the same genetic species, *homo sapiens*, but this seems implausible and not justified by Scheffler's analysis of the thought experiments (and the probably intuitive reactions thereto). It seems more plausible to suppose that the relevant general characteristic is simply that the future generations consist in beings living lives of rough *phenotypic* equivalence to those we now live (i.e. lives with similar cognitive and evaluative frameworks, similar projects, values and aspirations, willingness to preserve and build upon knowledge, art and other traditions).

The fact that it is this type of connection, and not the connection involved in biological/genetic parenthood, that matters, opens up the window to the argument being defended here because it is possible, as shall be argued below, that artificial offspring fit the bill just as well as natural offspring.

To briefly summarise, Scheffler maintains that the collective afterlife is an important sustainer of meaning in our lives because our understanding of what makes life valuable and meaningful depends on the satisfaction of *nonexperiential*, *conservative* and *altruistic* conditions of value for future generations. That is to say: we care about more than what we experience; we care about the preservation of things we value after we die; and we care about those things not just because they affect people with whom we have a biological and personal connection.

This suffices for the first step of the argument, but some further modifications of Scheffler's account are needed in order to ease the transition to the second step. Scheffler's thesis is purely concerned with the role of the collective afterlife in sustaining value and meaning. It is not concerned with other sources and sustainers of value and meaning. This limited focus is understandable: Scheffler is trying to defend a novel claim and it would distract from that novelty to focus on other factors that sustain the value and meaning in our lives.⁸ But it is important that those other factors are not ignored. Premise (4) of the argument being defended in this article depends on the comparative claim that creating artificial offspring is better than the currently available means of ensuring that there is a collective afterlife. To set-up this comparative claim, it is worth thinking about how the collective afterlife builds upon and relates to other conditions of meaning. To do this we need some sense of what these other conditions might be.

There are three general schools of thought about what it takes to live a meaningful and valuable life.⁹ They are:

Subjectivism: In order to live a meaningful life, an individual must attain or satisfy some subjective (i.e. mind-dependent) conditions, e.g. fulfil their desires, experience conscious pleasure, satisfy their interests etc.

Objectivism: In order to live a meaningful life, an individual must satisfy some objective (i.e. not mind-dependent)¹⁰ conditions, e.g. they must produce

⁸ Though it should be noted that some of his original critics focused on these other conditions in their responses, e.g. Wolf and Frankfurt in Scheffler (2013).

⁹ These schools of thought are discussed and defended in many sources, e.g. Smuts (2013), Metz (2013) and Campbell and Nyholm (2015). The latter argue that there are four main theories of meaning: (1) simple subjective; (2) simple objective; (3) aim-achievement and (4) fitting fulfillment. The classification system adopted here is a simplified version of this one insofar as it collapses the last two theories into one *hybridist* category of theories.

¹⁰ 'Objective' here does not, and should not, be understood to mean something like 'true in a scientific sense' or 'capable of being verified by a third party. This is sometimes the meaning that attaches to the term when people refer to 'objective' facts and the need for 'objectivity'. Here, 'objective' simply means that the conditions in question are not dependent on the mental states of the individual whose life they render more meaningful.

morally valuable outcomes, make some significant intellectual discovery, or produce aesthetically valuable art (Smuts 2013; Metz 2010).¹¹

Hybridism: In order to live a meaningful life, an individual must satisfy some combination of objective and subjective conditions, e.g. they must be consciously fulfilled by producing objectively valuable outcomes (Wolf 2010).

In addition to this general tripartite division between the different schools, there is a division between supernatural (predominantly theistic) and naturalistic schools of thought on the meaning of life (Weijers 2013). This division can be viewed as a dispute about the precise content of the subjective and/or objective conditions that need to be satisfied in order for life to have value and meaning. Thus, for example, a theist might argue that an individual must play a part in God's plan (an objective state of affairs) in order to live a meaningful life (Craig 2008). Contrariwise, a naturalist might argue that a life spent satisfying one's desires, or doing (wholly naturalistic) moral good would be enough.

Scheffler's afterlife thesis fits within this general taxonomy of views. It is hybridist and naturalistic. It is about the subjective fulfillment we receive from the satisfaction of an objective condition—namely: the existence of a collective afterlife. It is also wholly naturalistic: as mentioned above, Scheffler is very clear that his envisioned afterlife is not one that takes place in a supernatural heaven; it is the one that continues in our naturalistic universe.

But the point in describing this general taxonomy is not to classify Scheffler's thesis. The point is to highlight how it connects up to other theories about what it takes to add meaning to a life. Objective and hybridist theories point to the role of fulfilling and objectively valuable projects, such as the project of scientific discovery and exploration, the production of art, and the minimisation of moral suffering and wrongdoing, in sustaining and enhancing the meaning of our lives. This adds an intriguing complication to Scheffler's analysis. While the probable existence of a collective afterlife could be an important condition of meaning in our present lives, it could also be that there are better or worse collective afterlives (Johnston 2014; Kauppinen 2014). It could be that the better kinds of collective afterlife will be the ones in which other conditions of meaning are more likely to be satisfied. In particular, since the collective afterlife thesis is about the continued existence of beings sharing some of our evaluative beliefs and practices, it would seem like the better kinds of collective afterlife will be the ones in which fulfilling and objectively valuable projects are kept going and, if possible, improved upon. Those kinds of collective afterlife would ensure that we are sitting in the best possible chain of lives: one in which valuable projects are preserved and amplified.

This hints at another important condition for meaning in the collective afterlife, one that is missed in Scheffler's account. One thing that has been consistently thought to undermine meaning is *finitude* or *fragility* (Craig 2008; Di Muzio 2006; Weijers 2013). This is a charge supernaturalists repeatedly level against naturalists.

¹¹ Metz identifies these three things as the major sources of meaning in life, but does not defend a purely objectivist theory of meaning.

They argue that a major deficiency in the naturalistic worldview is that the lives we and others live can easily come to an end (see Di Muzio 2006 in particular on this point). We can cure someone of a life-threatening disease—thereby ensuring that we do some objective moral good—but this is a stop-gap measure: that person's life will eventually come to an end. We can also produce wonderful paintings and sculptures, but they will decay and crumble in the long-run.

There are famous critiques of the notion that finitude and fragility undermine meaning in life (Nagel 1971), but there are, equally, defenders of the naturalistic view that concede that such things do undermine meaning even if they do not completely rob us of it (Di Muzio 2006; Weijers 2013). This disagreement cannot be resolved here. Nevertheless, an assumption that will have repercussions for the argument in the next section will be made. It will be assumed that, even if finitude and fragility do not completely undermine meaning in life, it would, in general, be preferable if we could ensure a collective afterlife which is slightly more existentially robust than it currently is. Thus two modifications are being added to Scheffler's collective afterlife thesis. It is being claimed that the better kinds of collective afterlife will be the ones in which objectively valuable and fulfilling projects are continued *and* in which the pursuit and continuance of these projects is less fragile (i.e. not easily susceptible to corruption, loss, reversal and so on). These modifications are compelling as they take onboard certain criticisms of Scheffler's thesis. For instance, Mark Johnston (2014) has claimed that Scheffler's thesis in its original form has the unwelcome conclusion that life is a Ponzi scheme: we need an endless stream of future generations living lives of meaning and value to pay it down the line in order for our present generation to live lives of meaning and value. While one can disagree with aspects of this criticism—it seems unlikely, for example, that the problems faced by some distant future generation transfers all the way back to completely undermine the value of life in the present (Kauppinen 2014)—it does seem plausible to suppose that there should be a reasonable prospect of continued and indefinite survival and prosperity for those distant future generations.

To sum up so far: the first premise of the argument depends on the view that Scheffler is right: the existence of a collective afterlife (on this earth or elsewhere in the natural universe) is an important source and sustainer of meaning and value in our present lives because our conception of meaning and value has *nonexperiential*, *conservative* and *altruistic* dimensions to it. This core thesis had been modified by adding to it the claim that it would be better (best) if the beings in this collective afterlife continued (and helped to achieve) ongoing projects of objective value (e.g. moral, scientific or artistic), and if in doing so they were not readily susceptible to fragility and finitude, i.e. if they had a greater degree of existential robustness that would help to sustain them in the continuance of what we care about.

The Superiority of Artificial Offspring

We now come to the crux of the argument being defended in this article: premise (4). As you recall, that premise makes two sub-claims. The first is that the creation of artificial offspring is a credible pathway to the collective afterlife; the second is

that it may be a better pathway to the collective afterlife. Nothing elaborate will be said in support of the first sub-claim. A definition of what would count as artificial offspring has already been offered; what is meant by the collective afterlife thesis has also been clarified. It is relatively easy to see how the former (as defined) links up with the latter. To put it very simply, Scheffler's afterlife thesis claims that the value and meaning of our lives depends on the existence of future lives of 'rough equivalence' to our own. The 'rough equivalence' involves sharing our cognitive and evaluative frameworks, caring about similar projects and traditions. It seems plausible to suggest that sophisticated robotic offspring, with a close-to or slightly greater-than human level of artificial intelligence, would meet those criteria.

But there are several objections to this. Some may argue that the biological connection with human offspring is what matters. But, as has already been pointed out, this is implausible. Genetic and biological connection does not seem to be relevant when you take into consideration the thought experiments that support the collective afterlife thesis. What matters is some shared set of phenotypic characteristics. A more interesting variation on this objection would be to make the point that biological fragility is one of the qualities that must be preserved in the collective afterlife, and that this would be removed in the case of my hypothesised artificial offspring. There is something to be said about this, but we will return to it after we have considered why artificial offspring are better than biological offspring. We will be in a better position to process the argument at that stage.

The sentience of the offspring might be another point of contention. Some people might argue that the fact that biological offspring can experience things like pleasure or pain (or other more complex phenomena) is what matters, and they might be suspicious about the prospect of non-biological sentience. Sentience is quite probably a valuable property in future generations, but its importance does not block premise (4) from going through. In the first instance, artificial sentience is neither conceptually nor technologically impossible. In the second instance, it is argued here that we have to adopt a pragmatic, inferentialist approach to the sentience of all future generations. We can never know whether another human being is actually conscious, or shares our experiences of the world. We can only assume this by analogy with our own experience, or infer this from their outward behaviour. This is the case for our own offspring, and it is also going to be true for any artificial offspring: if they are like us and their outward behaviour suggests sentience, then this criterion will be met. Or, to put it another way, the possibility of non-sentient artificial offspring should cause no more angst than the possibility that all future non-artificial offspring are philosophical zombies. The same basic reasoning applies to anyone who insists that artificial offspring cannot sustain a meaningful collective afterlife on the grounds that they are not *persons* (in the philosophically rich sense of that term). There are competing accounts of the criteria that makes an entity a person—rationality, sense of self, verbal communication, consciousness etc.—but there is no reason to suppose that an artificial offspring (of the sort depicted in science fictional paradigms) would lack the behavioural manifestations of these criteria. Consequently, there is no reason to think they would be less satisfying bearers of our collective afterlife than human beings.

A related objection might appeal to the insights drawn from embodied cognition theory and argue that artificial offspring as they are being envisaged here are not technically feasible.¹² Embodied cognition theory highlights the various ways in which our biological nature shapes how we understand, interact and think about the world. The concern would be that, if embodied cognition theory is correct, the kinds of values we have, and projects we are interested in, are also likely to be shaped by our biological natures. If artificial offspring do not share our biological natures—if they are less prone to disease etc.—then what hope is there that they will form a collective afterlife that satisfies the conditions outlined above? There are three things to be said in response. First, it seems feasible for us to create embodied artificial agents that have similar embodied interactions with the world around them. These interactions may suffice to form a worldview broadly similar to our own. Second, even if you are skeptical of that claim, the kinds of values and projects that are important in the collective afterlife, as outlined so far, do not seem like ones that have much connection to our embodied, *biological* nature. A desire to pursue knowledge, the creation of art and the alleviation of moral suffering (to take the major examples) do not seem to require a biological constitution with its needs and susceptibility to disease. If anything, these projects only attain importance in human life once we have satisfied and moved beyond our basic biological needs. Third, even if one thinks there might still be something missing from these artificial offspring, one has to balance the losses they might experience (relative to us) against the gains. If they can pursue projects that we currently value but elude our grasp (as is argued below) that might compensate for their loss of other values that are unique to creatures with a biological constitution like ours.

Finally, another objection to the first part of premise (4) might be that my hypothesised artificial offspring are a long way off, technologically speaking, and hence it is difficult to see how they could sustain the meaning and value in our *present* lives. This might be a fair criticism, but two replies suggest themselves. First, developments in robotics and artificial intelligence have been proceeding apace in recent years, and given the widespread exponential growth in computing technologies (Kurzweil 2006; Agar 2015), the reality of full-blown artificial offspring might be closer than we realise. Second, for the argument to work, it is not absolutely essential for the proximate next generation to consist of some artificial offspring. In other words, I do not need to witness the creation of an artificial being before I die in order to ensure that I gain the benefits from this potential form of collective afterlife. All that matters is that humanity as a whole is committed to a project of creating such beings and this project has a high probability of success. There are some who will simply insist that we will never create artificial offspring of the sort hypothesized, and that this project has a low probability of success, but the evidence from technological progress seems to go against this belief.

This suggests that the first sub-claim within premise (4) is reasonably secure, which gives us *a* reason to think that the creation of artificial offspring might be worthwhile. Attention then turns to the second sub-claim, which is far more significant. Anyone who accepts the argument up until this point may respond: even

¹² I am indebted to an anonymous reviewer for drawing this objection to my attention.

if artificial offspring could provide us with a collective afterlife capable of sustaining meaning and value, why should we care? Aren't biological offspring a perfectly serviceable and more practically feasible method of ensuring that there is such an afterlife?

Concessions can be made to this point of view. The argument being defended here does not necessarily entail that we should pursue the creation of artificial offspring 'over and above' the creation of biological offspring. It is possible to have two parallel tracks toward the collective afterlife, each running simultaneously. That might be all for the better. Nevertheless, it is conceded that this is not that this argument is trying to establish. It is also trying to establish that there is something potentially superior about the artificial pathway to the collective afterlife. This requires a more complex argument that comes in two halves. The first half suggests that artificial offspring are better than biological offspring with respect to satisfying the broader set of conditions needed for sustaining meaning in life. The second half suggests that artificial offspring are better than biological offspring because they help us to avoid some of the tricky ethical issues involved in procreation. This means that there are no strong defeaters for the artificial pathway to the collective afterlife; and some also reasons to favour it.

Let's consider first how artificial offspring might be better in terms of satisfying the conditions for meaning in life. Some of this argument was presaged at the end of the previous section. As noted there, standard accounts of meaning in life appeal to the need to fulfill certain objective ends with our actions. As Metz puts it, great meaning in life seems to be tied to the pursuit of three ends in particular: the good, the true and the beautiful (2010). Our lives acquire great meaning to the extent that we conform our intellects to the pursuit and attainment of these ends. Obviously, this is something we can do in the course of our own lives, but projects directed toward the good, the true and the beautiful will also be something we wish to see preserved and continued in future generations.

Traditional biological offspring can, of course, allow for this. We all hope that our children and future generations will benefit from and continue our moral, intellectual and aesthetic projects, building and improving upon them as they go. And to some extent our hopes are likely to be realised (if the experience of past generations is anything to go by). But this is where the non-human-like aspects of artificial offspring can provide us with even more hope. These offspring will not suffer from the same fragilities and limitations as biological offspring. They could have enhanced minds that allow them to achieve greater scientific insights; they could be tireless in their pursuit of moral and intellectual excellence; they could be designed to be less vulnerable to weakness of the will; and less prone to distractions. In this way, we can be more hopeful that they will succeed in preserving and fulfilling our preferred moral, intellectual and aesthetic projects.¹³

In addition to this there is the possibility that artificial offspring can provide us with greater hope of overcoming finitude and achieving a collective afterlife of

¹³ It is not being assumed that they will have super-enhanced minds (i.e. be superintelligent). In fact, it is very important for the argument that they are not drastically different from us. Their abilities should be within or slightly beyond the normal human range. That way they can continue, sustain and improve upon our projects, without being so different that they no longer care about those project.

indefinite duration. As noted previously, the role of finitude in the meaning of life is somewhat contested: many have, historically, thought it to be essential; but some argue that it is not. This is partly due to the ascendancy of naturalistic theories of meaning that embrace the finality of death. Still, there are three reasons for thinking that overcoming finitude and achieving an indefinite collective afterlife could be relevant to the present argument. The first is that even among those who cling to a naturalistic account there is a belief that it would be better if we could overcome our finitude (Di Muzio 2006; Weijers 2013). The second is that there is a growing movement now arguing that we should incorporate philosophical uncertainty into our axiological projects (Lockhart 2000; Guerrero 2007; Moller 2011). Thus, even though we might be pretty sure that overcoming finitude is non-essential, we might still recognise the possibility of its importance, and so we should favour projects that try to address it, whilst still satisfying other conditions of meaning. In effect, we should *hedge our axiological bets*. People who embrace this view admit that not all risks of error should alter our axiological projects when tradeoffs are involved: some threshold of uncertainty would need to be crossed (Moller 2011). This threshold is not precisely defined, but it is submitted that it is crossed in the present case: there is sufficient uncertainty about the role of finitude in our theory of meaning, and no apparent cost from trying to overcome it, that it seems wise to hedge our axiological bet. The third reason is that the problem of finitude has already been identified by critics of Scheffler's afterlife thesis. As mentioned, Mark Johnston (2014) suggests that if we accept Scheffler's claim we must buy into a Ponzi scheme view of life. In other words, we must accept the need for an infinite stream of future generations paying it down the line to the present generations, despite knowing that there will not be an infinite stream. Johnston's Ponzi scheme argument might be flawed, but again this might be place in which we would wish hedge our bets: even if we don't need an infinite stream of future generations, it might be better if we could pursue a pathway to the collective afterlife that has a higher probability of providing it.

This leads naturally to the question: how might artificial offspring enable us to do this? Dan Weijers argument in favour of scientific optimism and techno-utopianism (2013) is instructive in this regard. Weijers's argument is that scientific and technological advancements contribute to a utopian view of human life because they offer up the hope that our actions in the present could have everlasting reach. But how can they do this? We know that our collective afterlife on earth is limited; that the earth will be destroyed by the explosion of our sun; and that future generations will, eventually, need to move beyond this solar system if they wish to continue in existence. The problem is that biological humans are not well-equipped for such a departure. We are too tethered to the environment here on earth: we have developed symbiotic and incompletely understood relationships with the local flora and fauna; we rely on food, water and a delicately balanced set of climactic conditions; we must breathe oxygen and avoid solar radiation; and we don't live long enough for extended space voyages. There are possibly ways that our biological descendants could overcome these limitations, but it seems like the easiest way, by far, would be to develop artificial offspring who are not so tethered to the conditions here on earth. They could live indefinitely, certainly long enough for interstellar travel; they would not need oxygen to breathe or water to drink; they would not be so vulnerable to

cosmic rays and radiation; they would not have to live in a symbiotic equilibrium with other evolved organisms. They would consequently be better equipped to continue the human project on other planets and in other solar systems. This would make the connection between our present lives and eternity a real possibility. The fact that they could do so would make an afterlife with artificial offspring better than one with biological offspring.

Someone might object to this on the grounds that it makes one fatally flawed assumption. It assumes that the universe itself is infinite in duration, and this is an assumption that might be rejected by some of the leading cosmological theories. But there are two responses. First, as Weijers points out, all existing theories on the fate of the cosmos are widely acknowledged to be incomplete and hence defeasible. There are theories that allow for an everlasting universe. Second, the argument being defended here does not necessarily require infinitude to be successful. The argument is a comparative one not an absolute one. It is claiming that artificial offspring are *better* than biological ones not that they are 'the best' (assuming that even makes sense). They could still be better in terms of overcoming *some* of the limitations and fragility of biological life, and offering up some hope of continuing collective afterlife until the end of the universe, even if they cannot provide hope for eternity.

So much for the first half of the defence of the claim that artificial offspring provide a better path to the collective afterlife. We turn now to the second half, which claims that artificial offspring are better than biological ones in avoiding the axiological and ethical pitfalls associated with human procreation. Despite the fact that having biological offspring is a natural and perpetual feature of human society, it is also widely acknowledged to carry with it a number of axiological and ethical risks (Overall 2012). For starters, sexual reproduction involves an interaction between at least two people.¹⁴ The ethical propriety of this interaction cannot simply be taken for granted.¹⁵ Likewise, at present, sexual reproduction involves an asymmetrical distribution of burdens as between the male and female. As Overall (2012) points out, it is the woman who bears the majority of that burden and this reality has frequently been neglected or improperly considered in human history. There have been recent advances in social attitudes towards women, and advances in artificial methods of fertilisation, but the asymmetrical distribution of reproductive burdens remains. The great advantage of artificial offspring is that they could be created without having to navigate our way through these ethical quandaries. They would not need to be conceived through sexual intercourse or through an agreement between two or more people. They would not need to be gestated by women. The burdens of procreation need no longer be distributed unequally across society. Now, to be clear, it is not being supposed that if we create artificial offspring, the future will be devoid of sexual relationships, biological offspring and gender based inequality. Again, the point is a comparative one: to the extent that artificial

¹⁴ The phrase 'at least' is used because reproduction involving IVF and/or surrogacy can involve more than two parties.

¹⁵ Consent and competence in child-bearing and rearing being the two major ethical issues involved in these cases.

offspring can help us to minimise or avoid the axiological risks associated with biological procreation, they have an advantage over biological offspring.¹⁶

In addition to this, artificial offspring may have the advantage with respect to the potential harms suffered by those living in the collective afterlife. Biological reproduction doesn't just carry risks for the parents of the child; it also carries risks for the child. The child may suffer from genetic or developmental disorders that render its life one of great suffering and pain. The child may not be loved or competently cared for. This can create repercussions for society down the line. Artificial offspring need not be vulnerable in the same way. They could avoid genetic and biological disorders. They would not necessarily need or require parenthood (though this may be a feature). They could be created fully-formed, so to speak, in a state of adulthood. This doesn't mean that they wouldn't have other problems, or vulnerabilities, but they would again have this advantage over biological offspring.

More controversially, artificial offspring could have the advantage over biological offspring with respect to the harm of coming into existence and the harm of death. The notion that life itself might be a great harm is controversial, as is the related claim that we should, consequently, not bring anyone into existence. Benatar is the most well-known defender of this view (2006). He rests his argument on two main claims. The first is an a priori claim based on the notion that there are asymmetrical harms and benefits associated with the act of procreation. In particular, that in bringing someone into existence you are responsible for all the bad things they suffer from but not for all the good things. The second claim is a posteriori and highlights the fact that human life is filled with more suffering than we typically realise. Most of the examples identified by Benatar have to do with suffering imposed by our biological limitations (cold, hunger, pain etc.) which we downplay because we either become desensitised to them or we compare them to other states in which they are corrected (i.e. states of warmth, satiation and pleasure).

Benatar's argument is certainly controversial, with many people resisting both its a priori and a posteriori components (Harman 2009; DeGrazia 2010; Overall 2012). But this is one more area in which we could once again acknowledge the possibility of philosophical error—i.e. the possibility that Benatar is right—and try to incorporate that possible error into our axiological projects. In doing so, the advantage that the project of creating artificial offspring may have over that of creating biological offspring can be seen. Many of the a posteriori harms identified by Benatar could be overcome in a suitably designed artificial being. They need not experience cold or hunger or other niggling complaints associated with biological fragility. They may not need to suffer from the raw experience of pain because they need not be wired to experience pain in the same way as we do. We need to experience pain so as to help avoid certain biological and existential threats: this has been hardwired into us through biological evolution. Artificial offspring would

¹⁶ It is not denied that there are certain joys and values associated with human procreation (beyond mere sexual pleasure). It is important to remember this article is not claiming we should shut down the biological pathway to a collective afterlife; it is merely claiming that the artificial pathway could be superior.

probably need to avoid such threats too, but they could have some internal warning system that alerts them to the existence of such threats without necessitating the raw experience of pain. Thus artificial offspring could avoid at least some of the harms of coming into existence.

The harm of death is also controversial. Many accept that death is a harm, with the most popular account of that harm being that death deprives you of the goods of continued existence. There are disagreements about this. The classic Epicurean position holds that death is not a harm because it is a *non-experiential* state (Warren 2006). But even Epicureans (among whose ranks the present author count themselves) concede that the process of dying can be harmful, even if the state of death is not. The relevance of this to the present dialectic is that some people may argue that there is something troubling about bringing into existence beings that will suffer from this great harm. It may not be enough to outweigh the benefits of coming into existence, but it is still something to consider. It would be better if we brought into existence beings that did not need to suffer this harm. Again, artificial offspring would seem to have the advantage over human offspring in this regard. They could potentially live indefinitely (with appropriate maintenance) and even if they are rendered obsolescent or have an equivalent of death, their process of dying need not be filled with the same kinds of suffering as our process currently is.¹⁷

In conclusion, there seem to be good reasons to endorse premise (4) of the argument. Artificial offspring provide a credible route to the collective afterlife, and the route they provide has an advantage over that provided by biological offspring with respect to fulfilling the objective ends needed to sustain meaning and value, and with respect to avoiding the axiological risks and harms associated with biological procreation and existence.

Objections and Replies

There are four additional objections to the argument that are worth considering before concluding. They each target different premises of the argument, suggesting that it either overlooks some crucial feature of what makes the collective afterlife a source and sustainer of meaning and value in this life, or avoids listing some cost associated with the creation of artificial offspring that obviates the benefits just mentioned.

The first objection is that in comparing artificial offspring to biological offspring a false dichotomy is assumed. The argument assumes that the choice is between traditional, biologically limited and fragile human beings, and more existentially robust artificial beings. But there is another possibility: the possibility of *enhanced biological offspring*, ones that no longer suffer from the sorts of limitations that put the biological offspring at a disadvantage. This is a real possibility, one that some people are eagerly advocating (Savulescu et al. 2011) but its possibility does not undermine the argument being defended here. For one thing, any sufficiently

¹⁷ The phrase 'currently is' is used because the prospect of universally available euthanising methods might reduce the harm for human beings.

enhanced biological offspring—one that overcomes the kinds of fragility mentioned in the previous section—would, it is argued, simply be equivalent to an artificial offspring: it would be a manufactured product of human invention and skill; its innate/natural biology would have to remarkably differ from that of any natural human. For another, there are moral risks associated with tampering with biological humans that would not apply to the creation of artificial beings from scratch. Agar highlights one of them. He notes how any attempt to overcome the limitations of human biology must grapple with the *integration problem* (Agar 2013, 48). That is to say, the problem that the enhancing technology must integrate with the pre-existing biological structure. This is a technologically and ethically fraught engineering project. As Bostrom and Sandberg point out (2008), the human body is the product of complex and delicately balanced evolutionary forces. Tinkering with it when it is improperly understood can often fail and backfire. Experimentation might be permissible if there was no real risk of harm to the entities being enhanced, but of course there is such a risk in the case of human beings. Developing an artificial being from scratch does not face the same engineering or ethical dilemmas: we are not tampering with a complex and incompletely understood evolutionary system.

The second objection is that in emphasising the fact that artificial offspring can overcome biological fragility and vulnerability the argument ignores the fact that those very qualities are essential to our conception of value and hence to our conception of why the collective afterlife is important. Something along these lines is suggested by people who analyse the unique good in parent–child relationships (Brighouse and Swift 2006). They point to the special vulnerability and dependency of the children as being central to that good. This is also suggested, more broadly, by Scheffler himself in his work on the collective afterlife. He claims that human vulnerability to existential threat is crucial to many of our value-laden concepts, so much so that a life without these vulnerabilities may not really be imaginable (2013, Chapter 3).

Two responses suggest themselves. First, it is not supposed here that the future artificial offspring would be immune from all existential vulnerability. They would still, presumably, be vulnerable to violent attack, accidents, lack of a power source, and possibly other non-biological threats. Also, they may have some initial level of vulnerability and dependency when they come into the world and try to learn about it (though this may not be essential). Second, and more importantly, it is possible to demonstrate that a future collective afterlife, consisting of such vulnerable artificial offspring, would sustain the same meaning and value as one consisting of biological offspring. We can demonstrate this by simply running Scheffler's thought experiments once more, only this time substituting (in our imaginations) artificial offspring, as they were earlier defined, for biological ones. Suppose that the next generation will consist entirely of artificial beings, sharing something like our cognitive and evaluative frameworks (e.g. Data-like beings with many similar hopes and aspirations to our own). Suppose further that this future generation will be wiped out by some global catastrophe (e.g. meteor strike, or global computer virus pandemic) thirty days after you die. Does this induce a similar level of disquiet and despair as Scheffler's original thought experiment? It is submitted that it does and

that this shows how we can be equally invested in the success of a collective afterlife consisting of artificial beings.

The third objection claims that in highlighting the potential advantages of artificial offspring over biological offspring, it was neglected to mention some of the potential costs. The obvious ones being the potential waste of valuable resources that might go into the project of developing such offspring, which might have been better spent alleviating present suffering or ensuring a better *human* future, or the potential existential threat that might be posed by smarter-than-human robots (Bostrom 2014; Armstrong 2014). Again, these are valid concerns, but they do not detract from the argument being made. With regards to the waste of resources, it is worth bearing in mind that resources are frequently being wasted on projects of limited moral worth (development of the latest social media app, or expenditure on the military in certain countries). Redirecting some of those resources to the project being proposing here might simply maintain an ethical status quo. But this is not all. If the argument presented above is correct, then the project being proposed has some genuine moral (axiological) worth: it could help to sustain the meaning and value of our present lives and give us more hope for the future. Those benefits should be weighed against the potential costs. This ties into a response to the existential threat that might be posed by such artificial beings. The argument does not claim that we should be prioritising this project over everything else, or that we should be pursuing it in a reckless and haphazard manner. It only claims that this is a project we have good reason to consider. Of course, we should ensure that other projects of value are not compromised; of course we should proceed with care and diligence. We already have a (biological) pathway to the collective afterlife, we can start gradually building the better (artificial) pathway now.

A fourth and (for now) final objection highlights another potential cost of pursuing the artificial pathway to the collective afterlife: increased social inequality or stratification. There are a couple of ways to make this objection. Some might argue that because artificial offspring would be created/manufactured beings, and not naturally conceived beings, they would be viewed as products and not as moral equals. Hence a collective afterlife with such beings would be one in which they are treated as inferior servants to their human creators. Alternatively, one could argue that because the artificial beings imagined in this paper are assumed 'superior' to humans in various ways, they might assume a superior status in society and foster inequality as a result. Either of these eventualities might cast doubt on the alleged advantages of artificial offspring over biological ones. There are several things to bear in mind when responding to this objection. First, note that the distinction between that which is created (and hence a product) and that which is naturally conceived is already increasingly blurred and is likely to be more so in the future. The blurring of this boundary doesn't seem to lend itself to an inferior status for the 'created' offspring. Nowadays, many adults do not just *happen* to have children. They *choose* to do so, frequently spending great sums of money on fertility treatments and surrogacy arrangements, and engaging in crude methods of genetic selection, in the process. In a very real sense, the children they breed are their 'creations' (born through conscious deliberation, intention and design). Indeed, popular objections to fertility treatments such as surrogacy and embryo selection

often point to the fact that the children produced through these arrangements are akin to products. This does not seem to have resulted in an inferior class of human beings. There is a good normative reason for this too. The moral status that we afford to one another should not be affected by our ‘ontological’ histories; it should be determined by the kinds of properties we exemplify in the present moment. If offspring are human-like in their phenotypic properties, then they probably won’t be (and shouldn’t be) treated as morally inferior.¹⁸ Second, to the extent that the artificial offspring might be viewed as (or view themselves as) morally superior and hence this might be catalyst for inequality, it should be noted that our current moral trajectory is shifting very much away from the belief that superior mental and physical capacity warrants unequal treatment. Look, for example, at the revolution in disability rights over the past half century. We are now beginning to embrace the notion that those with disabilities are *merely* different from the rest of us, not different in inferior ways. We are reducing the level of inequality and stratification experienced by disabled populations as a result (though we still have some distance to go). This recent experience should give us hope for a future in which we have artificial offspring who are better than us along certain physical and mental dimensions. They could live in harmony with our biological offspring, with both populations treating one another as merely different not badly/superiorly different. All that said, the concerns about future inequality should not be ignored. They are something to be factored into the mix when considering the wisdom of creating a collective afterlife of artificial offspring. Recall, from the outset, the goal has been to provide *a* reason (hopefully a strong one) for thinking that we should create such offspring, but not to provide an *all-things-considered* reason. Fears about inequality and other social costs may weigh sufficiently against the project. That is something we should decide, in our collective wisdom.

Conclusion

To sum up, this article has argued that the creation of artificial offspring—with the right mix of human and non-human qualities—might be better at sustaining the meaning and value of our present lives than a collective afterlife of ordinary biological offspring. There were two major components to this argument. The first was a modified defence of Scheffler’s collective afterlife thesis. According to Scheffler’s original thesis, the continued existence of human-like beings, in conditions of rough equivalence to our present lives, is an important source and sustainer of meaning in our lives. This thesis was modified by appealing to other objective conditions that might also be relevant to the sustenance of meaning. The second component of the argument was the claim that the development of artificial offspring provided both a viable pathway to the collective afterlife, and one that could be superior to the existing biological pathway. This was argued to be the case

¹⁸ This also highlights the flaw in thinking that simply because they are manufactured artificial beings cannot count as our ‘offspring’. If the manufactured nature of a being prevents it from being our offspring, then it would follow that children born through artificial selection or future genetic engineering are not our offspring (or are becoming less offspring-like all the time).

because artificial offspring might be better able to meet the criteria for meaning in the collective afterlife, and because their creation could avoid some of the axiological and ethical problems associated with human procreation.

It was mentioned at the outset that this was an unusual thesis. It may still feel a little unusual to the reader, but it is hoped that the resistance to it has been lessened and that the potential axiological benefits of artificial offspring are, if not compelling, at least more tempting.

Acknowledgements Funding was provided by Irish Research Council (New Horizons Grant).

References

- Agar, N. (2013). *Truly human enhancement*. Cambridge, MA: MIT Press.
- Agar, N. (2015). *The sceptical optimist*. Oxford: OUP.
- Armstrong, S. (2014). *Smarter than human*. Berkeley, CA: Machine Intelligence Research Institute.
- Benatar, D. (2006). *Better never to have been: The harm of coming into existence*. Oxford: OUP.
- Bostrom, N. (2014). *Superintelligence: Paths, strategies, dangers*. Oxford: OUP.
- Bostrom, N., & Sandberg, A. (2008). The wisdom of nature: An evolutionary heuristic for human enhancement. In N. Bostrom & J. Savulescu (Eds.), *Human enhancement*. Oxford: OUP.
- Brighouse, H., & Swift, A. (2006). Parents' rights and the value of the family. *Ethics*, 117, 80–108.
- Campbell, S., & Nyholm, S. (2015). Anti-meaning and why it matters. *Journal of the American Philosophical Association*, 1(4), 694–711.
- Craig, W. L. (2008). The absurdity of life without god. In W. L. Craig (Ed.), *Reasonable faith* (3rd ed.). Crossway: Wheaton.
- DeGrazia, D. (2010). Is it wrong to impose the harms of human life? A reply to Benatar. *Theoretical Medical Bioethics*, 31(4), 317–331.
- Di Muzio, G. (2006). Theism and the meaning of life. *Ars Disputandi*, 6, 128–139.
- Guerrero, A. (2007). Don't know, don't kill: Moral ignorance, culpability, and caution. *Philosophical Studies*, 136, 59–97.
- Harman, E. (2009). Critical study: David Benatar *Better Never to Have Been*. *Nous*, 43(4), 776–785.
- Johnston, M. (2014). Is life a ponzi scheme? *Boston Review*, 39, 42–48.
- Kauppinen, A. (2014). Why afterlifism isn't a Ponzi scheme. *PEA soup blog*. <http://peasoup.typepad.com/peasoup/2014/01/why-afterlifism-isnt-a-ponzi-scheme.html>.
- Kurzweil, R. (2006). *The singularity is near: When humans transcend biology*. London: Penguin.
- Lockhart, T. (2000). *Moral uncertainty and its consequences*. Oxford: OUP.
- Metz, T. (2010). The good, the true and the beautiful: Toward a unified account of great meaning in life. *Religious Studies*, 47(4), 389–409.
- Metz, T. (2013). *Meaning in life*. Oxford: OUP.
- Moller, D. (2011). Abortion and moral risk. *Philosophy*, 86, 425–443.
- Nagel, T. (1971). The absurd. *Journal of Philosophy*, 68, 716–727.
- Nozick, R. (1974). *Anarchy, state and utopia*. Harvard: Harvard University Press.
- Overall, C. (2012). *Why have children? The ethical debate*. Cambridge, MA: MIT Press.
- Savulescu, J., ter Meulen, R., & Kahane, G. (2011). *Enhancing human capacities*. Oxford: OUP.
- Scheffler, S. (2013). *Death and the afterlife*. Oxford: OUP.
- Smuts, A. (2013). The good cause account of meaning in life. *Southern Journal of Philosophy*, 51(4), 536–562.
- Warren, J. (2006). *Facing death: Epicurus and his critics*. Oxford: OUP.
- Weijers, D. (2013). Optimistic naturalism: Scientific advancement and the meaning of life. *Sophia*, 53(1), 1–18.
- Wolf, S. (2010). *Meaning in life and why it matters*. Princeton, NJ: Princeton University Press.