

Forthcoming in *Royal Institute of Philosophy Supplements*. This is the accepted manuscript version, which has not yet been copy-edited.

## **Meaning and Anti-Meaning in Life and What Happens After We Die**

*Sven Nyholm*

### **Abstract**

The absence of meaningfulness in life is meaninglessness. But what is the polar opposite of meaningfulness? In recent and ongoing work together with Stephen Campbell and Marcello di Paola respectively, I have explored what we dub ‘anti-meaning’: the negative counterpart of positive meaning in life. Here, I relate this idea of ‘anti-meaningful’ actions, activities, and projects to the topic of death, and in particular the deaths or suffering of those who will live after our own deaths. Connecting this idea of anti-meaning and what happens after our own deaths to recent work by Samuel Scheffler on what he calls ‘the collective afterlife’ and his four reasons to care about future generations, I argue that if we today make choices or have lifestyles that later lead to unnecessarily early deaths and otherwise avoidable suffering of people who will live after we have died, this robs our current choices and lifestyles of some of their meaning, perhaps even making them the opposite of meaningful in the long run.

On September 26 of 1983, a nuclear early-warning system in the Soviet Union mistakenly indicated that a missile had been launched from the United States, that it was heading in the direction of the Soviet Union, and that the missile had been followed by

five more. According to Soviet military protocol, it was the job of the duty officer of the command center to report the indications of the satellite warning system to officers higher up on the chain of command. The responsible lieutenant colonel of the air defense forces – Stanislav Petrov – decided to disobey his orders, however. He did so because he deemed this to be a false alarm. If he had relayed the message generated by the missile detection system in accordance with protocol, this might have triggered a large-scale retaliatory nuclear attack by the Soviet Union on the United States. A nuclear war might have broken out. Because he refused to convey the warning created by the faulty detection system, Stanislav Petrov is thought to have helped to prevent a nuclear war, with all of the devastation that it might have caused to large parts of the world. Petrov died in 2016 at the age of 77. During his lifetime, he was not reprimanded by the Soviet powers for what he had done, but nor was he rewarded for it. He retired early, and at one point suffered a nervous breakdown. During his lifetime Petrov did not gain much recognition for what he had done. In fact, it was mostly unknown to the rest of the world. Now that this has become more widely reported, however, we can all rejoice in what Petrov did for the world in disobeying his orders and, thereby, helping to prevent a nuclear war.<sup>1</sup> Even if it might not have led to happiness and public recognition for him during his lifetime, this can be seen as a very meaningful thing to have done in one's life.

Consider next a much more well-known person, who is still alive, but whose actions might also potentially benefit people living in the further future: the Swedish teenager Greta Thunberg, who for the last few years has worked tirelessly to raise awareness about the urgency of taking decisive action to try to prevent further human-

---

<sup>1</sup> Pavel Aksenov, 'Stanislav Petrov: The Man who May have Saved the World', *BBC News*, 26 September 2013: <https://www.bbc.com/news/world-europe-24280831> (accessed on July 30, 2021); Sewell Chan, 'Stanislav Petrov, Soviet Officer Who Helped Avert Nuclear War, Is Dead at 77', *New York Times*, 18 September 2017: <https://www.nytimes.com/2017/09/18/world/europe/stanislav-petrov-nuclear-war-dead.html> (accessed on July 30, 2021)

created damage to the climate. It is too early to tell whether Thunberg's actions and campaigning to fight climate change will help to turn the tide and thereby benefit future generations, who are currently at great risk of having to live in a world with a severely damaged climate. However, if Thunberg indeed manages to play a key role in helping to create a better world for people of future generations, as she might do, then this would certainly be another example of actions in the here and now that could be thought to be tremendously meaningful because of the good they do for future generations.

Others, in contrast, do things during their lifetimes that might seem to be the opposite of positively meaningful because of the harm or risk of harm that they bring about to people living after they are gone. Early during his presidency, for example, the former US president Donald Trump was all but outright challenging Kim Jong Un of North Korea to a nuclear war, when Trump was appearing in front of the UN's assembly and threatening the 'little rocket man' with 'fire and fury'. That could potentially have provoked a nuclear war and led to destruction on a massive scale, which might have led to suffering to many people who would live after Trump would no longer be around. Such actions – as well as other examples of actions and policies carried out by the Trump administration, such as actions denying the reality of climate change and seemingly threatening the future of US democracy – can seem like they are not simply meaningless as in being devoid of meaning. Rather, they can seem like actions that are the polar opposite of positively meaningful. In terms of the legacy that is left behind to those who will come after him, former president Trump's legacy will strike many of his critics as in no way being a positively meaningful one, but as something better described in the opposite terms.

These are three rather extreme cases. Most of us do not have opportunities to have such clear effects or possible impacts on those who will live after we are gone. But these

examples nevertheless help to illustrate the topic I will discuss in this paper: namely, the question of whether what happens after we are dead might help to make what we do in our lives more or less meaningful. I will be particularly interested in whether harm that we might do to future generations of people who will exist after we are dead might be thought to bring what I will call ‘anti-meaning’ (viz. the polar opposite of positive meaning) into our lives. The suggestion I will be discussing is that yes, indeed, what happens after we die – especially if it depends on our actions, projects, or lifestyles – might detract from the positive meaning of what we do and how we live our lives.

Discussions about meaning in life and its relation to death usually take a different perspective. They are usually about whether the fact that we will all die should be taken to threaten the possibility of living a meaningful life. Some philosophers – e.g., some transhumanists<sup>2</sup> – have argued that because we will ultimately all die, nothing we do in life can be truly meaningful, since it will all come to an end. Other philosophers – such as Leon Kass<sup>3</sup>, to pick just one example – have argued that the fact that we will die and that our time on Earth is limited is better viewed as something that can help to make what we do with our lives much more meaningful than it would be if our lives just went on and on and we never died. On this second way of seeing things, our mortality creates a certain urgency to everything we do, and helps to make each day we are alive and each project we undertake matter more. As I noted above, in contrast, I am here going to connect meaning in life and death in a very different way: I will discuss whether what happens after we are dead might affect how meaningful what we do during our lives can and should be thought to be. I will do so in a slightly gloomy way by focusing, as also

---

<sup>2</sup> See, for example, Max More, ‘Transhumanism: Towards a Futurist Philosophy’, *Extropy* 6 (1990), 6–13

<sup>3</sup> See Leon R. Kass, *Leading a Worthy Life: Finding Meaning in Modern Times* (New York: Encounter Books, 2017)

mentioned above, on how bad things that happen after we are dead might threaten the meaningfulness of what we do during our lives.

The discussion below divides into the following sections. First I will briefly say a little more about meaningfulness in life and the idea of ‘anti-meaning’, i.e. the idea of the polar opposite of positive meaning in life, a topic I have written about before together with Stephen Campbell and, more recently, Marcello di Paola (section 1).<sup>4</sup> Next I will bring up some ideas that might already have come to mind to some readers, namely, Samuel Scheffler’s work on what he calls the ‘afterlife conjecture’.<sup>5</sup> I will highlight some key features of Scheffler’s fascinating ideas about the importance of what happens after our deaths for what matters to us in our lives, and I will relate those ideas to some of the ideas my co-authors and I have discussed regarding what we call ‘anti-meaningful’ actions or activities (sections 2 and 3). This will lead me to my main claim, which is that our effects on people who live after we are gone can impact the meaningfulness of what we do during our lives, including in negative ways in case our activities negatively affect those who come after us. I will briefly suggest some examples, like issues to do with climate change, potentially devastating wars and other conflicts, as well as risks related to powerful AI technologies (section 4). Next, I will very briefly discuss some philosophical distinctions I think we should pay attention to when we reflect on these kinds of ideas (section 5), and then I will end with a concluding discussion (section 6).

## 1. The Notion of Anti-Meaningful Actions and Activities

---

<sup>4</sup> Stephen Campbell & Sven Nyholm, ‘Anti-Meaning and Why It Matters’, *Journal of the American Philosophical Association* 1 (2015): 694—711; Sven Nyholm & Stephen Campbell, ‘Meaning and Anti-Meaning in Life’ in Iddo Landau (ed.), *The Oxford Handbook of Meaning in Life* (in press); Marcello di Paola & Sven Nyholm, ‘Climate Change and Anti-Meaning’, under review.

<sup>5</sup> Samuel Scheffler, *Death and the Afterlife* (Oxford: Oxford University Press, 2011); Samuel Scheffler, *Why Worry about Future Generations*, (Oxford: Oxford University Press, 2018).

Shelly Kagan has an interesting article in which he discusses what he calls ‘ill-being’.<sup>6</sup> By this Kagan means the opposite of well-being. For example, if one accepts a hedonistic theory of well-being, on which well-being consists in experiencing pleasure or happiness, ill-being would amount to pain and suffering. On a preference satisfaction theory of well-being, ill-being would consist in preference frustration. In the same way, some authors writing about meaningfulness in life have found it useful to not only distinguish positive meaningfulness from meaninglessness (viz. the absence of meaningfulness), but to also distinguish meaningfulness from its polar opposite.

Different philosophers who have been interested in this idea have used different terms to pick out the same general notion. Typically, they use examples to make this idea intuitively plausible to readers. Thaddeus Metz, for example, talks about what he calls ‘anti-matter’. To make this idea seem intuitive, Metz asks his readers to contemplate the following:

Consider a life composed of actions such as killing one’s spouse for the insurance money or blowing up the Sphinx for fun. Does this life merely lack meaning or does it contain antimatter? Is it more akin to an impoverished life, one that lacks money, or to an unhappy life, one that contains misery?<sup>7</sup>

Metz expects us to think that such a life is more like the second of the two options: it is more like an unhappy life that contains misery than a life lacking in wealth.

---

<sup>6</sup> Shelly Kagan, ‘An Introduction to Ill-Being’, *Oxford Studies in Normative Ethics* 4 (2015), 261—288

<sup>7</sup> Thaddeus Metz, ‘Recent Work on the Meaning of Life’, *Ethics* 112 (2002), 781—814, 806—807

Similarly, Iddo Landau discusses the importance of ethics and moral considerations for meaning. According to Landau, if some activity is immoral, this not only means that it may lack positive meaning. It may be the ‘inverse’ of meaningful. Take, as an example, Adolf Hitler’s role in the mass murder of millions of people and other crimes against humanity during World War II. Landau writes:

Hitler’s life is not merely meaningless, a life in which meaning is absent, but instead a life in which the converse of meaning is present. To use an arithmetic analogy, the meaningfulness of lives such as Hitler had should not be evaluated at around zero or even as simply zero, but in negative numbers. The scale of meaning of life should be conceived as stretching into that negative sphere as well.<sup>8</sup>

Ingmar Persson and Julian Savulescu, in turn, use the phrase ‘bad meaning’ to refer to a similar idea.<sup>9</sup> They suggest that, in general, we should understand meaningfulness in a person’s life in terms of what the person intentionally does and what effects the person intentionally produces. If the effects of a person’s intentional actions are good, then this gives good meaning to the person’s life. In contrast, if what a person intentionally does has negative consequences, then this gives bad meaning to the person’s life.

Together with Stephen Campbell, and more recently also with Marcello di Paola, I have written about this topic as well.<sup>10</sup> We have used the term ‘anti-meaning’. Our main claim in these papers is that, as citizens of the modern world, it is very hard for most of

---

<sup>8</sup> Iddo Landau, ‘Immortality and the Meaning of Life’, *Journal of Value Inquiry*, **45** (2011), 309—17, 317

<sup>9</sup> Ingmar Persson & Julian Savulescu, ‘The Meaning of Life, Equality and Eternity’, *Journal of Ethics* **23** (2019), 223—238

<sup>10</sup> Campbell & Nyholm, ‘Anti-Meaning in Life and Why it Matters’, Nyholm & Campbell, ‘Meaning and Anti-Meaning in Life’, Di Paola & Nyholm, ‘Climate Change and Anti-Meaning’

us not to participate in activities that, in some fashion or other, adversely impact present or future people and/or serve to cement various forms of social injustice. This is not to say that ordinary people's lives are always or even typically anti-meaningful on the whole. Rather, the point is that various aspects of participating in modern society can serve to counteract, if not undermine, the meaningfulness of some aspects of our lives or lifestyles.

## 2. Scheffler on the Afterlife

In my discussion about anti-meaning in this paper, I am particularly interested in whether what happens after we die might make some of the things we do while we are alive less meaningful or perhaps even to some extent anti-meaningful. Readers familiar with Samuel Scheffler's recent work<sup>11</sup> might be reminded of it here, and I am indeed also strongly influenced by Scheffler in my thinking about this. Scheffler uses the term 'the afterlife' to refer to what happens to people who live after we ourselves die. And he uses the following types of scenarios to motivate what he calls the 'afterlife conjecture' regarding the importance of the afterlife:

*The Asteroid Scenario:* a gigantic asteroid is traveling through space towards the planet Earth. It is a little bit like in the motion picture *Armageddon*, in which an asteroid is traveling towards Earth and the only way to save the planet is to blow up the asteroid. In the film, Bruce Willis portrays a hero who is able to destroy the asteroid and save everyone. In Scheffler's scenario, there is an asteroid approaching, but no Bruce Willis-type of hero, and we know that everyone will

---

<sup>11</sup> Scheffler, *Death and the Afterlife*, Scheffler, *Why Worry about Future Generations*

soon die because the world will be hit by the enormous asteroid. So, there is no future for human life on Earth.

*The infertility scenario:* in a second, less violent scenario, there is no giant asteroid that is about to dramatically collide with Earth and thereby immediately kill everyone. Rather, just like in the novelist P.D. James's book *The Children of Men*, everyone on Earth has been unable to conceive any children for the last 25 years. Otherwise, everyone is apparently healthy and appears to be living out a normal lifespan. Since nobody is having any new children, the world's population will eventually die, and there will be no future people coming after us. We are the last generation – not because of a nuclear war or because of an explosion killing everyone in one go – but rather because we simply are no longer having any children.

Scheffler expects us to react with horror and dismay, or at least sadness and a sense of resignation, when we contemplate these kinds of scenarios. His 'afterlife conjecture' is that what happens to other people after we ourselves die matters much more to us than we may know. Interestingly, that is the sort of language Scheffler uses when he talks about this idea. He talks about how what happens to people after we die 'matters' to us more than we perhaps realize. He says that we 'care' about this. And he says that 'our capacity to value things' here and now depends on our thinking that there will be an afterlife after we are gone. We can accept, Scheffler thinks, that we ourselves will die. But we would find it much harder to accept a situation where there is no afterlife, viz. when there is nobody who lives after us in the future.

Strikingly, Scheffler does not use the language of what is ‘meaningful’ in life when he talks about this topic. Instead, he talks primarily, as just noted, about what is valuable to us, what we care about, or what matters to us. But it seems like it would have been a natural choice of terminology to put things in terms of what is meaningful to us, or what helps to make our current activities meaningful to us. Perhaps this is just a matter of what terminology Scheffler is comfortable with or prefers to use, and what fits best with the terminology of a lot of his previous work, for instance on what is involved in valuing something.<sup>12</sup> At any rate, it appears very intuitive to interpret what Scheffler is talking about in terms of what helps to give meaning to what we do in our lives.

When it comes to what Scheffler is aiming to do, he presents his project as being one of finding reasons to care about future generations that are (i) non-experiential, (ii) non-utilitarian, and (iii) personal in nature. Since we care about what happens to people in the future after we die, and we will have no experiences of the future after we die, our reasons for caring must be of a non-experiential kind, Scheffler argues. He is also trying, he says, to present an alternative to the sort of broadly utilitarian or altruistic reasons for caring about future generations that Derek Parfit discusses in his monumental last part of *Reasons and Persons*.<sup>13</sup> The reasons Parfit discusses to care about future generations are related to how we can benefit future generations – but Scheffler is interested in different kinds of reasons, of a non-utilitarian character. Lastly, Scheffler is trying to steer away from impersonal concern with the future and towards more personal reasons for concern, which have to do with things that we care deeply about in our own lives. Another thing that could be noted here is that ‘reasons’ here seems to refer to what Bernard Williams

---

<sup>12</sup> E.g. Samuel Scheffler, *Equality and Tradition: Questions of Value in Moral and Political Theory*, (Oxford: Oxford University Press, 2010)

<sup>13</sup> Derek Parfit, *Reasons and Persons*, (Oxford: Clarendon Press, 1984)

has famously called ‘internal reasons’.<sup>14</sup> That is to say, the reasons Scheffler is concerned with are reasons that derive from what we care deeply about, as opposed to ‘external reasons’ that are not grounded in our own internal cares and concerns. Scheffler does not himself explicitly pledge allegiance to a Williams-inspired internalism about reasons. But he nevertheless bases his whole line of argument on attempts to show what we care about deep down, whether or not it is immediately obvious to us. That is the hallmark of an internal conception of reasons for action of the sort Williams has articulated.

In his second book on this topic, Scheffler goes on to articulate four such reasons to ‘worry about future generations’.<sup>15</sup> I will not discuss them in great detail here, but just quickly explain what they are. They are:

1: *Reasons of interest*: many of our projects that we take an interest in, Scheffler thinks, assume that others will carry on these projects after we are gone. If we, for example, are doing research about how to cure cancer (to use one of Scheffler’s examples), or if we are building bridges or other things it takes a long time to complete, then part of what we implicitly think gives value to these projects is that others will carry on these projects or benefit from what we do in the future, including after we are gone.

2: *Love of humanity*: When we are sad or dismayed about the prospect of their being no future generations, this shows, Scheffler thinks, that we have a deep ‘love of humanity’. We love humanity and therefore we would find it very sad and unfortunate if humanity would die out soon and there would be no humans

---

<sup>14</sup> Bernard Williams, ‘Internal and External Reasons’, in Ross Harrison (ed.), *Rational Action* (Cambridge: Cambridge University Press, 1979), 101—113

<sup>15</sup> Scheffler, *Why Worry about Future Generations*

shortly after we are dead. It is comforting to us, Scheffler thinks, that people will still be there after we are gone, because we harbor a deep love of humanity.

3: *Reasons of value*: It matters to us, Scheffler thinks, that there are people around in the future who can value the sorts of things that we value. There can be no value without valuers, as Scheffler sees things. And so, in order for the future to be a future of value, there need to be human beings – or other valuing beings – around who can value things and thereby carry on our traditions of valuing certain things or kinds of things.

4: *Reasons of reciprocity*: Lastly, since our lives derive some of their value from the existence of an afterlife in Scheffler's sense, it is not only the case that future generations can benefit from what we do for them. We can and do also benefit from their existence – e.g. they can carry on projects that we have an interest in and they can continue valuing things we find important. Therefore, it is mutually beneficial for us and future generations that we exist and that they are able to exist as well.

Like I said, I am not going to discuss in any great detail the reasons named by Scheffler for why we worry about future generations. Hopefully the short summaries of the reasons Scheffler articulates that are provided above are clear enough for the reader to get a sense of the sorts of concerns that Scheffler points towards. What I want to do here rather than discuss further Scheffler's own suggested reasons is to suggest that we should add a fifth general reason for caring about what happens to future generations. We should add what might be called a *reason having to do with meaningfulness*. That is to say, my suggestion

is that since what happens to other people after we are already dead might help to make our own lives more or less meaningful. Thus, a further reason to care about future generations – in addition to those formulated by Scheffler – is a reason related to how future generations can help to make our own lives more or less meaningful. Why? In short, this is because meaningfulness in life has to do, among other things, with how what we do fits into a larger picture. In other words, as I see things, things we do in life get their meaning, in part, from the larger context(s) of which they are part. This can be understood in different ways. For example, Antti Kauppinen has discussed a ‘narrative’ theory of meaningfulness in some of his work, on which meaning in life partly depends on the overall narrative of our lives.<sup>16</sup> On another view – defended by Ben Bramble – meaning in life has to do with whether we live in ways that produce good overall consequences.<sup>17</sup> I do not here mean to necessarily commit myself to one or the other of these views, but I think they are both on to something true about how it makes sense to think about meaningfulness: namely, to repeat, that it has to do with, among other things, how what we do fits into a larger picture.<sup>18</sup>

To briefly summarize the main claims of this section: firstly, although Scheffler himself uses other terminology when he discusses how it matters to us what happens to other people after we are gone, I suggest that what Scheffler is talking about can very plausibly be translated into claims about what makes a difference to how meaningful our lives are. Secondly, in addition to the sorts of personal reasons that Scheffler articulates for caring about what happens to future generations – reasons of interest, love of

---

<sup>16</sup> E.g. Antti Kauppinen, ‘Meaningfulness and Time’, *Philosophy and Phenomenological Research* **84** (2012), 345—377

<sup>17</sup> Ben Bramble, ‘Consequentialism about Meaning in Life’, *Utilitas* **27** (2015), 445—459

<sup>18</sup> In general I adhere to Susan Wolf’s ‘hybrid’ theory of meaning in life, whereby meaning in life is achieved when we are passionate about things that have a non-subjective value. This view has both a ‘subjective’ and ‘objective’ component. The latter can plausibly be related to the idea of doing things that are part of something ‘bigger than us’, as Wolf and many others put things. See Susan Wolf *Meaning in Life and Why it Matters* (Princeton: Princeton University Press, 2010).

humanity, reasons of value, and reasons of reciprocity – we should also include reasons of meaningfulness in our accounts of what different kinds of reasons we have for caring about people in the future.

### 3. Anti-Meaning and Activities that Harm – or Risk Harm to – Future Generations

Let's return again to the two scenarios from Scheffler's discussion that I outlined above: the asteroid scenario and the infertility scenario. In the former, there will be no people in the near future because an asteroid is about to collide with Earth and kill everyone. In the latter, there will be no people in the near future because all people have become infertile through no fault of our own. As noted above, Scheffler thinks that such scenarios pose a threat to our capacities to find value in the activities we engage in now, the idea being that our capacity to find value in our activities importantly depends on our assumption that others will come after us. If we follow the suggestion I made above and we think about these scenarios and Scheffler's 'afterlife conjecture' in terms of the concept of meaningfulness, the category related to meaningful that suggests itself is that of meaninglessness. That is, the way that Scheffler describes the reactions he thinks most people would have – and would justifiably have – to the scenarios he discusses seems to suggest that he is talking about reactions that indicate that we would view many of our activities as losing their meaning or becoming less meaningful if there is no 'afterlife' in Scheffler's sense.

Notably, the scenarios Scheffler primarily discusses are ones in which it is not human agency that lies behind the discontinuation of human life on Earth. In one scenario, there will be no more human life because an asteroid is about to kill everyone.

In the other, there will be no more human life because people have become infertile through no fault of their own. In neither scenario is it human agency – viz. anything that we have done – that has brought about the bad outcome that Scheffler thinks we will view as sad or disheartening. It is of great interest in this context, however, to also consider scenarios in which the reason why the future will have certain bad features is because of human agency. I therefore think that we should also consider scenarios such as these two broad kinds:

*Scenarios involving what Ingmar Persson and Julian Savulescu call ‘ultimate harm’<sup>19</sup> that is caused by what people do who live now: there is something that we do in the present that prevents the future possibility of human life on Earth (or any other form of life that we find highly important).*

*Scenarios that involve what we might call ‘irreparable harm’ that is caused by what people do who live now: there is something that we now do that makes life for people living after us very unpleasant, risky, or bad in other highly significant ways.*

Again, what makes these two kinds of scenarios radically different from the ones Scheffler focuses on is that in these two broad kinds of scenarios, the reason why there is either ultimate harm (no possibility of life on Earth) or irreparable harm (greatly decreased quality of life) is that people who live now act in some way, have certain forms of lifestyles, or in other ways cause these outcomes through our human agency.

---

<sup>19</sup> See Ingmar Persson & Julian Savulescu, *Unfit for the Future* (Oxford: Oxford University Press, 2012)

We can here imagine various different kinds of ways in which these kinds of scenarios could come about. The literature about so-called existential risks is a great source of possible ways in which these scenarios could come about.<sup>20</sup> For example, we can think of human-created climate change and excessive resource-depletion caused by the ways in which we live in the modern world. We can imagine cases in which we elect leaders who start or provoke nuclear wars. Or we can imagine scenarios in which powerful artificial intelligence (AI) systems or other technologies get out of hand and cause great damage.

Some of these scenarios might also be related to each other. For example, Aimee van Wynsberghe, Mark Coeckelbergh, and others argue that many current AI technologies are highly environmentally unsustainable by having a very big carbon footprint.<sup>21</sup> Another comment about risk of ultimate or irreparable harm caused by powerful AI that we may lose control over: one possible scenario is the sort discussed by Nick Bostrom in his book *Superintelligence*<sup>22</sup>, i.e. where AI systems become much more intelligent than human beings but where the goals they have been given are most easily achieved by them in ways that pose existential risks to human beings. However, this is not the only kind of out-of-control AI systems that we need to worry about and that could do great damage to us now or to people living after we are gone. So long as an AI system gets out of control and is powerful enough, it might cause a great deal of damage, even if the system is by no means super-intelligent. Consider again, for example, the nuclear attack detection system that generated the false alarm discussed in the Stanislav Petrov

---

<sup>20</sup> See, for instance, Toby Ord, *The Precipice: Existential Risk and the Future of Humanity* (London: Bloomsbury, 2020)

<sup>21</sup> See Aimee van Wynsberghe, 'Sustainable AI: AI for Sustainability and the Sustainability of AI', *AI and Ethics* (published online ahead of print, 2021) & Mark Coeckelbergh, *Green Leviathan or the Poetics of Political Liberty: Navigating Freedom in the Age of Climate Change and Artificial Intelligence*, (London: Routledge, 2021)

<sup>22</sup> Nick Bostrom, *Superintelligence* (Oxford: Oxford University Press, 2014)

example in the introduction. That was by no means a super-intelligent AI system. However, if Petrov had followed the recommendation from the detection system and warned his superiors of an on-coming (but actually non-existent) American nuclear attack on the Soviet Union, this might potentially have created a devastating nuclear war. The philosophically inclined computer scientist Roman Yampolskiy describes many such potential AI risk scenarios in his work that do not involve any superintelligence.<sup>23</sup>

If these kinds of scenarios occur – i.e. we do certain things now that make things extremely bad for people living after we are gone, or that make it impossible for there to be any more people after we are gone – what categories related to meaningfulness then suggest themselves? Here I want to submit that it is not only meaninglessness that seems to be the right category to think in terms of. That is, our doing things that cause ultimate or irreparable harm to future generations does not only seem to be devoid of positive meaning. Rather, the category of ‘anti-meaning’ seems to suggest itself as a better category to use here. In other words, if people living after we are gone will have it very bad or will not be able to live at all because of what we do now, then what we do now can appear to become the opposite of meaningful for this reason, at least when considered in light of these future effects of what we do now. In other words, the bad effects our current lifestyles, projects or choices might have on other people who will or would live after we die can add a dimension of anti-meaning to the things we do now while we are alive. When we put the effects of what we do now into the larger perspective of how it will affect those who come after us, this could give what we do a very positive meaning, but it could also give a negative meaning – or anti-meaning – to what we do now.

---

<sup>23</sup> See, for example, Roman Yampolskiy, *Artificial Intelligence Safety and Security* (Boca Raton: CRC Press 2018)

#### 4. Relevant Distinctions Related to Potentially Anti-Meaningful Aspects of Our Lifestyles and Projects

When we think about whether what happens after we die – and especially when we think about what happens after we die because of what we do now – there are some philosophical distinctions that might help to make it easier to specify what our conclusions about this issue should be when it comes to meaningfulness and anti-meaningfulness in life. I will not discuss all possible distinctions that might be helpful here, but simply focus on a sub-set I find particularly relevant.

Firstly, Dale Jamieson and Marcello di Paola draw an interesting distinction between what they call the ‘episodic’ and ‘systemic’ life of our actions and activities.<sup>24</sup> The episodic aspects of our actions are those that are most immediately obvious to us, and the most immediate effects of our actions. For example, when driving a car, what is obvious to us might be the fact that it gets us quickly to the destination we want to go to. But the act of driving cars is part of a practice that pollutes the climate and has various other effects that are problematic in nature, potentially both for current and future generations. Some acts are in themselves perfectly harmless – e.g. taking a long warm shower – but might be part of a system that uses up resources in an unsustainable way – e.g. lots of people taking lots of long warm showers.

This distinction is not exactly the same as but similar to the distinction between the short-term effects and the long-term effects of our actions, lifestyles, or projects. Again, the short-term effects may be relatively harmless. But the long-term effects of

---

<sup>24</sup> Dale Jamieson & Marcello di Paola, ‘Political Theory for the Anthropocene’, in David J. Held and P. Maffettone (eds.) *Global Political Theory* (Cambridge: Polity Press, 2016), 254—280

having lots of people engage in some lifestyle, use some technology, or having certain projects might have very bad effects, both for the people currently living and potentially also for future generations not yet living.

This, in turn, can be related to the distinction between foreseen (or perhaps unforeseen), but definitely unintended or at least not directly aimed-at consequences or side effects of our activities, lifestyles, or projects. Insofar as we live our lives in ways – or have projects – that might make things worse for people living after we are gone, this is highly likely to be unintended or at the very least not a directly aimed-at effect of how we conduct ourselves. Of course, there are some people who do wish to harm current and future people. For example, there are doomsday cults with leaders who devise plans for how to bring about the end of the world. But most people – the vast majority of people – are not like that. When we do things that bring harm to others (whether people in the present or people in the future), these are typically unintended side effects of what we are doing.

This, to bring up one last distinction, relates to risks associated with what we are doing, as opposed to the potential benefits or goods we are seeking to bring about. When we harm other people (present people or future people), this is often because we were doing something risky while pursuing some potential benefit or good that itself might not be harmful to others or that carries with it risks of harm to others as a side effect of achieving the good or benefit for ourselves.

With these distinctions drawn and concepts highlighted, I can now restate the point I was making towards the end of the previous section in somewhat more precise terms. My claim is that if our current ways of living/our projects have side effects whereby bad outcomes occur after our deaths (e.g. devastating climate change or AI disasters), then this can make our ways of living/projects less meaningful, to the point

that they might partly become the opposite of meaningful ('anti-meaningful'), at least to some extent. Some activities (e.g. trying to develop certain AI technologies) might be meaningful in the short term (because of their potential to do good), but anti-meaningful in the long run (e.g. because they are highly unsustainable and risk leading to devastating crises in the further future). By imposing great risks on people who will live after we are dead, we are in effect risking to make our own activities less meaningful.

In suggesting this claim, I am assuming the following three things about meaningfulness in general. Firstly, I am assuming, as noted above, that meaningfulness – as well as anti-meaningfulness – has to do with how our activities fit into a larger picture. Secondly, I am assuming that whether our actions turn out to be meaningful or anti-meaningful depends on what actual effects our activities end up having. Thirdly, I am taking it that having good intentions is not enough, so to speak, in order for our actions and activities to end up being meaningful and avoid being anti-meaningful. Kantians are right, as I see things, that whether we have a good will and are good people depends on what our intentions are and what we are trying to do. If we act with the best of intentions and with good motives, then we are good people with a good will even if things do not turn out in a good way. But when it comes to meaningfulness, I am taking it, things are different. When it comes to this issue, whether what we are doing and how we are living our lives can be seen as positively meaningful or ends up being partly anti-meaningful does depend on how things ultimately turn out.

## 5. Concluding Discussion

Above, I started with the example of Stanislav Petrov and how he might have ‘saved the world’ when he disobeyed Soviet protocol by refusing to relay the information from the satellite detection system that seemed to indicate that the United States had launched a nuclear attack on the Soviet Union. This, I suggested, can be seen as a tremendously meaningful thing to have done, since so many people potentially benefited from this action, since this might very well have prevented an all-out nuclear war. Let us suppose now, instead, that Petrov had followed his orders, that this might have triggered a nuclear war, and that this would have led to damage on a massive scale for the people who lived after Petrov and those around him were gone. And let us consider the people who had worked to develop the faulty nuclear attack detection system. Their intentions might have been good. But they would have developed a technology that led to suffering on a large scale both during their lifetimes and after their lifetimes, since a large-scale nuclear war is likely to cause damage that several generations might suffer from. In this counterfactual version of the example, what the engineers who developed this technology did, and what Petrov in this version of the story did when he passed on the information generated by the technology, can be seen as being the opposite of positively meaningful. It can be seen as anti-meaningful.

The worry I have been trying to formulate in this paper is that we – or some of us – might potentially be doing similar things. People developing powerful AI technologies might potentially be developing technologies that people will lose control over and that might lead to damage and destruction on a large scale. This might not be because the AI technologies are ‘super-intelligent’, as noted above, but simply because they are highly unsustainable from an environmental point of view. Or it might be because they lead people to act in destructive ways, like in the example above with the counterfactual story about how things might have gone in 1983 in the Stanislav Petrov case. To use a less

dramatic example: the fact that we are using up the world's resources at an alarming rate, and that we are doing so while polluting the climate, might also lead to very harmful effects for future generations, who will then need to suffer greatly after we are dead because of how we lived our lives while we were alive. This can be seen as detracting from the positive meaning of how we live our lives in the modern world, and as introducing – at least to some extent – the polar opposite of positive meaning into our lives in the modern world.

This can seem more likely to be the case when it comes to people who have a lot of influence over others, and who can therefore do more harm to current and future generations. To go back to another one of the opening examples: somebody like the former US president Donald Trump might have succeeded in convincing lots of people that climate change is 'fake news'. While the actions of somebody like Greta Thunberg might help to inspire many people to live in more sustainable ways, the actions of somebody like Trump might discourage many people from trying to live their lives in more sustainable ways. The actions of a Greta Thunberg can seem highly meaningful in a positive sense with respect to how they might impact people who live after she is no longer around, while the actions of a Donald Trump are likely to end up being the opposite of positively meaningful – or anti-meaningful – in the long run, because of the harm they might cause to people who will live after Trump is no longer alive. Most of us do not have the kind of influence that Thunberg and Trump have. So, we might not have opportunities to do meaningful things during our lives whose positive meaningfulness depends on great benefits enjoyed by those who are here after we are gone nor opportunities to do things that are strongly anti-meaningful because of great harms to those who will live after we are gone. However, by participating in the modern world – with all that this means in terms of depleting resources and causing climate change

through the systemic effects of our actions and activities – we all face the risk of introducing anti-meaningfulness or negative meaning into our lives because of the ways in which we live our lives now will affect those who will live after we are no longer around.<sup>25</sup>

Let us now also briefly return to the relation between what I have been suggesting and the ideas in Samuel Scheffler's fascinating books about what he calls the 'afterlife' in which he discusses 'why to worry about future generations'. How does what I have suggested above contrast and compare with what Scheffler puts forward in his work on this topic?

Firstly, as noted above, Scheffler does not put his claims in terms of what is meaningful or anti-meaningful. Rather, Scheffler is discussing what we 'care' about, what 'matters' to us, and what we 'value'. Of course, this does not mean that there is any necessarily very sharp distinction to be drawn between what Scheffler is interested in and what I have been interested in above. What we care about, what matters to us, and what we value are very closely tied to what appears meaningful to us or, more generally, to meaningfulness in life.

Secondly, I noted above that when Scheffler articulates reasons to care about future generations, he seems to be trying to identify what Bernard Williams calls 'internal reasons' to care about future generations. That is, Scheffler is articulating reasons that are closely – if not essentially – tied to our cares and concerns. This is contrasted with 'external reasons', which are reasons that we might be thought to have for doing or caring

---

<sup>25</sup> Some readers might be willing to grant that actions and activities that might harm future generations can be viewed as potentially being anti-meaningful to some extent in theory, but wonder why exactly this should be taken to matter in practice. In other words, why care about anti-meaning? I do not have enough space to discuss that important issue here, but I discuss it in some detail in Campbell & Nyholm, 'Anti-Meaning and Why It Matters', Nyholm & Campbell 'Meaning and Anti-Meaning in Life', and Di Paola & Nyholm 'Climate Change and Anti-Meaning'.

about things, which are wholly independent of what we actually do care about or would care about in certain counterfactual circumstances. It might be asked whether what I have been talking about – viz. reasons to care about future generations related to meaningfulness in our own lives – should also be viewed as a form of ‘internal reasons’ tied to what we already care about, or whether they can also be interpreted as ‘external reasons’ that we have whether or not we do care. Regarding this issue, I intend my discussion above to be neutral with respect to whether we should think of reasons as being ‘internal’ or ‘external’ in Williams’ senses. That is to say, while Scheffler seems, at least as I am reading him, to be particularly interested in identifying internal reasons to worry about future generations, the claims I have formulated above are not meant to necessarily be of that nature. I want to side-step the issue of how we should understand normative reasons. The only thing I need for my purposes is the premise that we can make value judgments about whether the effects of our current actions and activities on future generations are good or bad. How exactly this is to be interpreted from a more meta-philosophical point of view I wish to leave open here.

A third comment about the relation between what I have been saying above and what Scheffler says in his work: when I suggested above that in addition to the four reasons to care about future generations that Scheffler discusses, we should also add a fifth reason related to meaningfulness, I did not mean to suggest that such a further reason would be completely independent of or different from the four reasons to worry about future generations that Scheffler articulates. As a reminder, Scheffler’s reasons are reasons of interest, love of humanity, reasons of value, and reasons of reciprocity. These can in different ways be seen as being related to how meaningful our actions and activities are. And so the idea I am suggesting – viz. that what happens to future generations because of what we do matters to how meaningful our actions and activities

in the here and now are – can be seen as a form of or alternative formulation of one or more of the reasons Scheffler has articulated. Or it might be seen as a type of corollary of Scheffler’s conclusions about what reasons we have to care about what happens to future generations. The statement that our actions and activities might have good or bad effects on people who live after we are gone and that this might make those actions and activities more or less meaningful – and that this gives us a reason to care about what happens to future generations – fits neatly with the reasons to worry about future generations that Scheffler articulates, and is an alternative way of explaining the importance of the kinds of reasons that he puts forward.

A fourth and last comment about the relation between Scheffler’s discussion and mine: as I noted above, but as bears repeating, what Scheffler is interested in is simply the possibility that there will be no future generations soon after we are dead and the question of what this means for us. To bring this out, Scheffler focuses particularly on scenarios where it is not our fault that there will be no future generations. In contrast, I have been particularly interested in scenarios in which what we do here and now significantly impacts the lives of those who will live after we are gone. I have joined Scheffler in partly being interested in scenarios in which there are no more future people shortly after we are dead (e.g. because we start a massive nuclear war), but there too I have been primarily interested in such cases where this depends on what we do. And I have also been interested in cases in which there will be people living after we are dead but where the quality of life of those people has been significantly impacted – perhaps in a bad way – because of how we live our lives here and now. Scheffler is more narrowly interested in whether the mere existence or non-existence of future generations should be taken to matter to us now. My concern has been primarily with the impact of our actions and activities on people who live after we are gone – both in terms of whether

our actions might make it impossible for there to be many more future generations and in terms of whether we might have good or bad effects on the lives of people who are able to live after we are gone. This is a key difference between my focus above and Scheffler's focus in his work.

Let me lastly also note that I do not mean to suggest that meaning and anti-meaning in life only or even primarily depends on what happens after we are dead. Just like what we do while we are alive can be very meaningful because of what happens while we are still alive, what we do while we are alive can be, in certain respects, anti-meaningful because of what happens while we are still alive. None of what I have discussed above has been meant to suggest otherwise. I am in agreement with authors like Thaddeus Metz, Iddo Landau, and others that destructive and immoral actions and activities, for example, are among the things that might be anti-meaningful because of the negative value of these actions and activities in the here and now. Although my focus in this paper has been on how what happens after our deaths might have positive or negative impacts on how meaningful our activities turn out to be in the long run, this is not to deny that we can and should also concern ourselves with how meaningful our activities are in the short run. Both issues are very important. But the one I have focused on deserves more attention, which is why I have highlighted it in this paper.<sup>26</sup>

*Utrecht University*

*Email: [s.r.nyholm@uu.nl](mailto:s.r.nyholm@uu.nl)*

---

<sup>26</sup> For helpful feedback on this material, many thanks to Michael Hauskeller and the participants of his "Meaning in Life and the Knowledge of Death" *Royal Institute of Philosophy* 2020 conference, especially Frances Kamm and Antti Kauppinen. My work on this article is part of the research program Ethics of Socially Disruptive Technologies, which is funded through the Gravitation program of the Dutch Ministry of Education, Culture, and Science and the Netherlands Organization for Scientific Research (NWO grant number 024.004.031).