

Value in Very Long Lives

Preston Greene
Philosophy Programme, NTU Singapore
pgreene@ntu.edu.sg

Epigraph

We shall not cease from exploration, and the end of all our exploring will be to arrive where we started and know the place for the first time. (T. S. Eliot, *Little Gidding*)

Abstract

As things currently stand, our deaths are unavoidable and our lifespans short. It might be thought that these qualities leave room for improvement. According to a prominent line of argument in philosophy, however, this thought is mistaken. Against the idea that a longer life would be better, it is claimed that negative psychological states, such as boredom, would be unavoidable if our lives were significantly longer. Against the idea that a deathless life would be better, it is claimed that such a life would be lacking important sources of value, because death is a precondition for many of our valuing attitudes. I argue that these problems are avoided by very long (and potentially infinite) lives that incorporate fading memory, limited ignorance of future events, and temporal scarcity. I conclude that very long lives are, in principle, desirable, and that death does not play an essential role in our valuing attitudes.

Keywords

bioethics - eternal life - life extension - meaning of life - value theory

(1) Introduction

If we are lucky, we live for about a hundred years. Our lives progress through stages: infancy, childhood, adolescence, adulthood, and old age. We forget many things, and may be completely ignorant of the details of the early stages of our lives. We have some expectations of what our future will be like, though this varies. We often devote ourselves to a career — such as that of a philosophy professor — throughout our adulthood, which ends in retirement during old age. Regarding the future, the one thing we are nearly certain of is that one day we will die. This is human life as it currently stands.

Now imagine an alternative. You know you have lived for a very long time, though you are not certain how long. Your memory is as limited as it is in the actual world: you remember your life back to a point, though the details fade as you go further back in time. You know that you have 50 years remaining on your current contract as a philosophy professor, and you expect to do something else after that, though you are ignorant of exactly what. It seems possible that one day you will die, though of this you are not certain. You are certain that, given the limitations of your memory, one day you will no longer be able to remember the details of the way you now live, or the people you now care about, just as you currently cannot remember the details of how you lived so many years ago.

This is a vision of a very long (and potentially infinite) life that incorporates fading memory, limited ignorance of future events, and temporal scarcity. The first thesis of this essay is that the combination of these elements — or ones of a similar nature — can make such a life desirable. The second thesis is that death plays a merely contingent role in our valuing attitudes, and can be replaced without a perceived loss of value.

Why think that a very long life is undesirable or lacking in important sources of value? Several philosophers have raised concerns about lives that are very long, and some have concluded that living for a very long time would be bad, and that this is so simply due to the nature of very long life. We can distinguish between two types of concerns about very long lives. On the one hand there are *psychological-state problems*, which focus on the supposed bad effects of very long life on an agent's subjective psychological states. On the other there are *value problems*, which focus on the supposed bad effects of very long life on the value of an agent's activities, regardless of the agent's state of mind. Below I consider each type of problem in turn.

(2) Psychological-State Problems

Bernard Williams (1973) argues that a very long life would result in a problematic situation similar to that of Elina Makropoulos (or “EM”) — the 342-year-old title character of Karl Capek's play “The Makropoulos Affair.” Williams sometimes characterizes EM's situation as involving “boredom” (e.g., p. 90), and also sometimes as involving a lack of “categorical desires” (100). Categorical desires are contrasted with what Williams calls “conditional desires,” which are desires that only apply on the

condition that the agent survives. For example, you might desire to have a dental procedure performed on the condition that you are to continue living, but you do not desire to continue living so that you can have the procedure performed. The desire to have the procedure performed is therefore a conditional desire. Categorical desires, on the other hand, are not conditional on an agent's survival. It is for the satisfaction of these desires, Williams explains, that we may desire to continue living.

On an everyday understanding of “boredom,” it is clearly wrong to equate boredom with a lack of categorical desires. Often we are bored not because we lack categorical desires, but rather because there is nothing we desire to do that we are presently able to do (consider, for example, the plight of a prisoner in solitary confinement). The phrase “bored with life” may come closer to how Williams means to be describing EM's situation, and a plausible interpretation of “bored with life” may reference a lack of categorical desires.¹ In any event, Williams' view on the connection between boredom and categorical desire seems to be complex (if not overly complex²), and it is partly for this reason that I will formulate a response to Williams' concerns that does not focus on the avoidance of any particular psychological state, but rather the avoidance of negative psychological states generally (in Section 2.1).

Why does Williams believe that boredom would be inevitable in a very long life? Williams puts particular emphasis on constancy of character as a necessary condition for personal persistence over time. He identifies one's character with the set of experiences that have made up one's life and the set of experiences that one is disposed to seek out. Having a character, for Williams, will require both of these sets to be finite. This leads to the idea that it is a “fantasy” for an agent experiencing an infinite life to avoid boredom while retaining her character through time. As Williams (94) states, “One thing the fantasy has to ignore is the connexion, both as cause and as consequence, between having one range of

¹ Though, as Steele (1976, 425) and Wisnewski (2005) point out, a lack of categorical desires may be a temporary problem. So what Williams seems to require is something closer to “permanently bored with life.” Bortolotti and Nagasawa (2009) provide a detailed discussion of different types of boredom and how they relate to Williams' argument.

² Williams (91) states, for example, “In EM's case, her boredom and distance from life both kill desire and consist in the death of it.”

experiences rather than another, wishing to engage in one sort of thing rather than another, and having a character.” Since the range of experiences compatible with an agent's character is finite, if an agent retains her character through time she will eventually encounter all of these experiences, and boredom will ensue. (Note that this problem does not essentially concern infinite lives, as it would also apply to any finite life that is long enough to exhaust all of the experiences compatible with one's character.)

Can this problem be solved by long-life models in which an agent lives through a series of disjoint stages, without retaining memories of previous stages? These models fail to satisfy Williams' second condition for desirable long life. He states, “The second important condition is that the state in which I survive should be one which, to me looking forward, will be adequately related, in the life it presents, to those aims which I now have in wanting to survive at all” (91). Williams believes the memory-separated-stages approach does not satisfy this condition, since the memory separation between stages makes it impossible to establish an adequate connection between one's current desires and their satisfaction in the next stage. As Williams (92) notes, a stronger claim may be made by those who accept a psychological approach to personal identity that focuses on a memory criterion for persistence. According to a memory criterion approach, such a situation would not amount to personal persistence over time. However, Williams holds that the lack of connection to an agent's aims, on its own, is a sufficient reason for rejecting these models.

(2.1) A General Worry Regarding Supposed Psychological-State Problems

When it comes to psychological-state problems, Williams seems to have succeeded in identifying the most pressing concerns for those living a very long life. Being bored, lacking desires, or, more artistically speaking, having “too much of oneself” (Williams, 100) — these all seem to be of the best candidates for inescapable psychological-state problems in the context of very long life (note the contrast in plausibility with the claim that *anxiety*, *anger*, or lack of *belief* are inescapable). As Shelly Kagan (2012, 243) writes, “Essentially, the problem with immortality seems to be one of inevitable boredom. The problem is tedium.” Richard Taylor (2000, 331), in one of the most popular discussions of the meaning of life, is of the same opinion. He imagines an eternal being whose greatest desire is to raise a temple: “Let us suppose he succeeded in this, that after ages of dreadful toil, all directed at this final result, he did at last complete

his temple, such that now he could say his work was done, and he could rest and forever enjoy the result. Now what? What picture now presents itself to our minds? It is precisely the picture of infinite boredom!”

To what degree should we rely on our imaginative abilities in attempting to settle this debate? As Moore (2006) and Burley (2008, 81–2) have pointed out, arguments for and against necessary boredom in the context of an infinite life place heavy demands on, and perhaps unjustified confidence in, our imagination. If very long life is unimaginable, then we should not expect appeals to imagination to result in accurate affective forecasting. Consider Nagel (1986, 224)’s insistence that, unlike Williams, he cannot imagine ever becoming bored with life, and his potential diagnosis: “Can it be that [Williams] is more easily bored than I?” In light of this, Moore (2006, 314) suggests that whether one sides with Nagel or Williams reflects one’s “temperament” more than one’s evidence. I am inclined toward the idea that these criticisms reveal serious flaws in much of the literature concerning the desirability of very long life. Consider Temkin (2008, 193)’s disarmingly honest reflections on this debate: “Unfortunately, like everyone else, my views on this topic are based not on *experience*, but on mere speculation and imagination. Clearly, then, anything I, or anyone else, writes on this topic should be taken with a large grain of salt.”

Here is the approach I intend to pursue. I will first give, without appeal to imagination, a general defense of the idea that *no psychological-state problem* is a necessary consequence of very long life. I will then point out some particular ways in which the problems that Williams raises can be avoided, which will further support the general argument. Finally, I will discuss how this idea is consistent even with strict theories of personal persistence through time, such as the one Williams endorses.³

I note first that perhaps Williams is correct in thinking that human psychology is such that it would necessarily evolve into bad psychological states when left to its natural processes over a very long period of time. I think that evidence for or against this assumption is sparse, but will grant it for now. An

³ Metz (2002, 791) gives the interesting suggestion that boredom need not destroy the meaningfulness of one’s life, even if it does make life undesirable. Since I intend to argue that boredom is not a necessary feature of very long life, I will ignore this distinction for simplicity.

interesting question remains: in what ways could we have control over our psychological states, while retaining our identity over time?

Consider how things currently stand with the treatment of anxiety disorders. Someone suffering from generalized anxiety may experience chronic and exaggerated worry over the events in their life. From a certain perspective, the development of generalized anxiety can be seen as a consequence of the “natural” progression of one's psychological states in conjunction with a certain combination of life experiences. An effective form of treatment for this disorder is medication, which involves the adjustment of certain physical properties of one's brain. While there are costs associated with most forms of treatment (in the case of medication, most of the costs are contingent on current limitations in medical technology), losing one's personal identity is not considered one of them.

The situation is similar with mood disorders such as depression. (Depression offers a particularly apt example for our current purposes, since it is often associated — either as cause or effect — with what might rightly be called a lack of “categorical desire.”) As with anxiety, standard treatments for depression involve manipulation of the physical properties of one's brain. And as with the treatment of anxiety, the successful application of this treatment is not thought to result in a loss of identity.

How far can these sorts of treatments go? In thinking about this question, let us first assume physicalism, which is a very popular view. Applied to the philosophy of mind, physicalism is often defined as the view that mental properties supervene on physical properties. Mental supervenience does not entail the claim that the mental *reduces to* the physical, but it is compatible with that claim. If the mental reduces to the physical, then we can talk of the “reduction bases” of mental states — such as that of boredom — in terms of physical states. If so, then control over physical states implies control over mental states. This view is compatible with physicalist views that accept multiple-realizability, such as functionalism, since functionalists can accept token identities. However, without the reduction claim, the supervenience claim alone is slightly too weak to establish that control over physical states implies control over mental states, since supervenience only establishes that there cannot be a mental difference without a physical difference (and not that physical differences necessarily produce mental differences). Nevertheless, even nonreductive physicalists accept “mind-brain correlations.” The existence of such

correlations would also ensure that control over physical states implies control over mental states; if we can realize all possible physical states then we can realize all possible mental states as well.

For the purposes of determining the desirability of very long life, it is best to assume that any physical properties in the supervenience base of mental states are adjustable. To deny this would put the critic of very long life in the awkward position of claiming that we will never be able to adjust certain physical properties, regardless of advancements in future technology. This would be problematic for two reasons. First, it is hard to see how we could be in a position to know that certain future technologies, though nomologically possible, are unlikely to be developed. Second, what we are primarily interested in is whether a very long life would be desirable, in principle. A satisfactory answer to this question would not appeal to the idea that certain life-aiding technologies are unlikely to be developed.

From the physicalist's perspective, then, it becomes clear what a critic of very long life who focuses on psychological-state problems must accomplish. She must provide an argument with the conclusion that there is a necessary connection between 1) the preservation of our personal identity, 2) the passing of time, and 3) the existence of the problematic base properties. In other words, the argument must conclude: *The existence of these negative base properties is a necessary condition for personal persistence over a period of time.* I do not see a successful argument to this conclusion in Williams, or anywhere else. This is so even given the questionable assumption with which we started (viz., that natural human psychological processes tend to produce negative states if left unchecked).

One implausible route to this conclusion is the claim that interference with our natural psychological processes destroys personal persistence. This claim is not only inconsistent with accepted practices for the diagnosis and treatment of psychiatric disorders, it is also inconsistent with the major theories of personal persistence (see Section 2.3).

(2.2) A Response to Williams' Concerns

A more limited claim that is perhaps adequate for Williams' purposes is the following: it is essential to one's character that one become bored in response to repeated experiences. The truth of this claim might establish that boredom is inevitable given a lifespan long enough to ensure that the experiences

“compatible with one's character” are repeated. This, combined with Williams strict requirement on personal persistence — that one's character remains fixed through time — might seem to be sufficient.

Before getting to my main objection, let me note two initial worries for this type of claim. First, notice that the claim cannot be that it is essential to one's character that one would tend to become bored by repeated experiences *without adjustment from an external source*. This claim is too weak, since we are assuming the availability of technology that can keep us unbored. Rather, what is required is the implausibly strong claim that it is essential to one's character that one *actually* experience boredom in response to repeated events.

Second, people might differ in the degree to which repeated-experience boredom is a part of their character. In fact, people seem to actually differ in this regard. While some people eschew routine and constantly seek out new experiences, others appear, at first glance at least, to be content in following a strict daily routine. Imagine someone, for instance, who enjoys waking at dawn for a morning walk, spending a couple hours working in their garden during the day, and then enjoying dinner with a glass of wine from their balcony overlooking the sunset. It is not clear that this repeated pattern of experience must eventually become boring, and it is even less clear that the actual experience of boredom under such conditions must be an essential part of one's character. If this is so, then even the implausible-seeming claim that Williams requires results only in the conclusion that *for people of a certain character* very long life would inevitably become boring.⁴

Now for the main objection: for any psychological-state problem supposedly caused by the repetition of experiences, the key factor is not how repetitive an agent's experiences actually are, but rather how repetitive an agent's experiences *appear to be*. For example, imagine yourself at a restaurant you find appealing. Even if you have already eaten at this restaurant at this particular table and ordered this particular meal, you may still find the experience fresh and exciting if you have no memory of the prior event. What makes our experiences appear repetitive is our memory of their prior occurrences; lacking

⁴ Cf. Temkin, 2008, 203: “Some people bore easily, and others do not.” Fischer, (1994, 264–5) gives a vivid description of the type of character that is unlikely to become bored by repeated experience.

memory of prior occurrences allows us, seemingly, “to arrive where we started and know the place for the first time.”⁵

Surprisingly, Williams seems to assume that living a very long life entails an increase in one's memory capacities. Here is his description of a world in which everyone has an indefinite lifespan: “It would be a world of Bourbons, learning nothing and forgetting nothing, and it is unclear how much could ever happen” (90). Williams considers the possibility of incomplete memory only in terms of total memory wipes occurring at the transitions between discrete stages (and he rejects this possibility for the reasons discussed above).⁶ He may be making this assumption because he has a particular conception of the nature of eternal beings in mind,⁷ or he may be assuming that the technology that allows us to live indefinitely will also allow us to attain greater memory capacities. In any event, there is no reason to suppose that living longer entails that one's memory capacities increase.⁸ In fact, the most promising model of a desirable very long life, I suggest, is one in which our memories are neither perfect nor periodically wiped, but rather of a capacity that is adequate to our needs: relatively accurate regarding the near past but slowly fading over time. As long as the variability of experienced events outruns memory capacity, it is possible that a life of even infinite length may appear to the agent, at each moment in time, to involve *no repetition of experience whatsoever*.

⁵ Though I understand that this is likely not what T. S. Eliot meant in the context of *Little Gidding*.

⁶ Perrett (1986, 224–7)'s criticisms of immortal life operate within this same problematic framework.

⁷ This is not an uncommon feature in discussions of immortality. Consider, for example, Scheffler (2013, 97)'s assertion: “I do not think it is a trivial assumption ... that immortal beings would continue to be susceptible to pleasure and pain.” This statement can only make sense, I believe, given the assumption that the brains of beings who live forever would have to function in a way drastically different from our own.

⁸ As an associate editor of *Journal of Moral Philosophy* points out, there is an independent problem with Williams' claim if “forgetting nothing” is interpreted literally. Given plausible assumptions, the physical realization of infinite memory capacities requires the brain to be infinitely large.

An agent with a fading memory like ours might thus live a life filled with repeated experiences while avoiding the negative psychological states that repetition is supposed to cause, and all the while exhibit constancy of character. Indeed, our memory capacities as they currently stand seem appropriately suited to the task of making a very long life appear non-repetitive. If future people decide to leave their natural memory capacities unaltered, it might turn out to be a good idea to make sure that sequences of close experiential repetition are kept apart by at least a few hundred years.

(2.3) Theories of Personal Persistence

At this point it may be worthwhile to consider the ways in which theories of personal persistence are compatible with these strategies for avoiding psychological-state problems. We have, in fact, been assuming a very strict necessary condition for personal persistence — that one retains the same character over time — in order to engage with Williams' argument. If we relax this assumption, then it becomes even clearer that maintaining personal persistence is not a barrier to avoiding psychological-state problems throughout a very long life.

Bodily criteria of personal persistence are the easiest to accommodate, since they do not place any restrictions on the sorts of alterations that can occur to an agent's psychology over time. Therefore, there is no barrier to the avoidance of psychological-state problems once we assume the existence of technology capable of adjusting mental properties. The more relevant worries about very long life on bodily-continuity views concern how long our physical bodies can persist.

Psychological approaches that focus on a memory criterion of personal persistence provide more of a challenge. The worst-case scenario is a *direct memory criterion*, which requires a direct memory connection in order to establish identity. A classic problem with this type of view is that it seems to violate the transitivity of identity: *C* might recall experiences had by *B*, and *B* might have recalled experiences had by *A*, even though *C* cannot recall experiences had by *A*. This motivates memory criteria that require only *indirect* memory connections (or memory *continuity*) to establish identity: *C* is the same person as *A* if *C* can recall experiences had by *B* and *B* could recall experiences had by *A*. This is a welcome criterion for an advocate of the fading-memory strategy, since, ideally, one's memory would

fade over time in a way that preserves an indirect connection all the way back to the beginning of one's personhood.

However, it may be possible to salvage the direct memory criterion from identity intransitivity by utilizing contemporary theories of object persistence. According to “four dimensionalism,” objects persist through time like they extend through space. Objects persisting through time have “temporal parts” in the same way that objects extending through space have spatial parts. Applied to persons, these parts can be thought of as “person stages.” The question of personal persistence then becomes the question of what sort of relationship must hold between person stages in order for the stages to be of the same person, and establishing this relationship amounts to establishing a criterion of personal persistence.

Within four-dimensionalist theory, there is a disagreement between “space-time worm” theorists and “stage” theorists. Space-time worm theorists, such as Lewis (1976), identify persons with maximal aggregates of suitably-related person stages. Lewis avoids intransitivity by claiming that these aggregates can sometimes overlap. When a direct memory criterion is applied to the case above, Lewis' theory holds that *C* and *B* are person stages of one continuant person while *B* and *A* are person stages of another continuant person. In his discussion of longevity [29–31], Lewis claims that a case of someone like Methuselah — which ostensibly describes one person living for a very long time — might actually describe many people with overlapping stages. According to Lewis, each maximal aggregate of suitably-related “Methuselah” person stages represents a person.

Stage theorists, such as Sider (1996), identify persons with person stages, and analyze talk of persistence over time with an intransitive temporal counterpart theory. When a direct memory criterion is combined with stage theory, we get the result that *C* is a counterpart of *B* and *B* is a counterpart of *A*, but *C* is not a counterpart of *A*. Sider (2001, 204–5) claims that the case of Methuselah features counterpart intransitivities, and not overlap.

It may be possible, then, to make sense of the direct memory criterion, and such a criterion would rule out the fading-memory strategy for avoiding repetition-induced psychological-state problems, since as an agent's memory fades she eventually reaches a point at which she no longer qualifies as the same person. Even so, in order to demonstrate the existence of unavoidable psychological-state problems, this criterion must be combined with Williams' contention that consistency of character is a

necessary condition for personal persistence, as well as the problematic assumption that it is essential to one's character that one actually become bored in response to repetition. By itself, the direct memory criterion does not entail these additional necessary conditions, and thus allows for personal persistence in situations in which an agent avoids unpleasant psychological states through the use of brain-altering technologies.

Moreover, theorists who adopt psychological approaches to personal identity, including Lewis and Sider, endorse a broader notion of psychological connectedness than that of the memory criterion. If memory-connectedness is not a necessary condition for psychological connectedness, then there is no in-principle objection to the use of a fading-memory strategy to avoid psychological-state problems.

Furthermore, there are many theorists who accept broad-continuity over broad-connectedness.⁹ Whether a broad-continuity theory is compatible with strategies that involve brain modification depends on the finer details of the theory and the sort of modification required. However, all the strategies that we have discussed allow for perfect memory connectedness between adjacent stages, and it is implausible that the required brain modifications would necessarily be so extreme as to block personal persistence in spite of perfect memory continuity. This is because even broad-continuity theories of personal persistence agree that memory continuity is among the most important criteria for establishing identity through time.

(3) Value Problems

So far we have been considering arguments against the desirability of very long life that focus on the inescapability of psychological-state problems. A different type of problem concerns the value of a very long life. As with psychological-state problems, I will first present a general defense of the idea that no value problem is a necessary consequence of very long life. I will then point out some ways that particular value problems can be avoided.

(3.1) A Basic Alternative to Short Life

⁹ Examples include Shoemaker (1984) and Unger (1990).

Consider the following “basic” alternative to short life:

You live a life just like the one you will actually lead up until the point of death. However, rather than dying, you retire to a “pleasure machine,” where you experience a permanent (or very-long-but-finite) state of euphoria.

The example is interesting because it is identical to what critics of long life grant is valuable — the short lives we currently lead — in every way except for what happens at the end: instead of retiring to a very long period of nothingness you retire to a very long period of pleasure. Of how much value is this period of pleasure? According to a simple hedonistic theory, this period is immensely valuable. However, even according to theories which deny that simple euphoria is valuable, the basic alternative still seems to be a valuable life in so far as the beginning part is valuable (it is rarely claimed that pleasure is in itself a bad thing, or — more to the point — that experiencing pleasure is worse than experiencing nothing at all). In order to deny that the basic alternative represents a valuable life, one must pursue the seemingly-difficult path of showing how pleasure at the end *destroys* the value of one's activities at the beginning.¹⁰ If it does

¹⁰ It is possible to accept that the basic alternative represents a valuable life but to settle for the weaker claim that *at some point* more time alive must stop adding to the value of one's life (for example, the point at which the agent enters the machine). It is outside the scope of this essay to refute this claim. However, even so, the basic alternative would establish that very long lives do not create value *problems*, in the sense that they are of less value than shorter versions. Once this has been established, we can turn to the question of whether there is an upper limit on valuable additional life. I suspect that the upper-limit view is often generated by ill-founded concerns over psychological-state problems (Section 2) or a lack of temporal scarcity (Section 3.2). But there are exceptions. Consider Bramble (2015, 1078–80)'s view that repeated pleasures do not contribute to a person's lifetime welfare. Such a view might imply that there is an upper limit on the amount of welfare that the pleasure machine can provide. However, it does not imply that continued existence in the machine *lowers* overall welfare. Something similar can be said of Blumenfeld (2009, 371–7)'s “philosophy of life,” which assigns central importance to “real novelty.”

not, then we have an example of an infinite (or long-but-finite) life that retains the value of our finite lives, since the first part is valuable, and the life is infinite (or long-but-finite).

The critic's path is actually more difficult than it might first seem. Not only must she hold that pleasure destroys the value of one's activities at the beginning, but also that it does so in a way that the cessation of experiences (i.e., death) does not. Otherwise, the conclusion of her argument will support nihilism about the value of our lives (which is a view outside the scope of this essay).¹¹

A more plausible (and more likely) response to this case would be to deny that existence in a state of general euphoria would qualify as a "life" at all. Since existence in such a state is compatible with theories of personal persistence, this response requires that we divorce the concept of living a life from that of a person undergoing a succession of experiences, and that we adopt a stricter definition of the end of life than "the cessation of experiences." Even so, the example would still show that a very long "succession of experiences" is valuable, even if it does not show that a very long "life" is valuable.

However, in order to accommodate stricter definitions of "a life," we can adjust the basic alternative to involve an "experience machine" similar to that described by Robert Nozick (1974). Instead of dying, you are put into an experience machine (again, either permanently or for a very-long-but-finite amount of time) that provides you with a pleasure-inducing simulated life. Since we have already dismissed the idea that psychological-state problems are unavoidable over the course of a very long life,

When real novelty is exhausted, Blumenfeld argues, it would not be unreasonable for a person to choose to stop living. Even so, choosing to continue on would not destroy the value of what has come before. A possible view that does deliver the result that pleasure at the end destroys value at the beginning, which I do not discuss here, is that valuable lives necessarily exhibit a certain type of "shape" or "narrative structure," and that these shapes or structures necessarily include temporal finitude. For criticism of this view, see Fischer 2005, 279–80 and 2006, 397–402.

¹¹ Arguments in favor of nihilism that focus on the inescapability of death are endorsed or discussed by writers such as Craig (1994, 73), Taylor (2000, 332), and Tolstoy (1987, 31). Critics of very long life must make sure that their premises do not support this general nihilistic conclusion if their problem is to concern *long* life, rather than life in general.

we can assume that the machine really does work as advertised (indeed, perhaps the machine utilizes some of the strategies for avoiding psychological-state problems discussed in Section 2).

It is worth noting that Nozick's point in discussing the experience machine is not so extreme as that life in the machine is worse than death, or even that life in the machine is valueless. Rather, his conclusion is that a life outside the machine is better than a life inside, even if life inside the machine promises more pleasure. For our purposes here, we can accept this conclusion, and even the stronger conclusion that life in the experience machine is valueless. A critic of very long life must make the additional claim that entering the machine destroys the value of one's prior activities.

(3.2) Temporal Scarcity

It might well be granted that there is no reason that long-but-finite lives must be valueless. Nevertheless, one might hold that the situation is different with infinite lives, because infinite lives lack an end. In his recent book, Samuel Scheffler (2013) argues that the lack of an end makes valuing impossible in the context of an infinite life, because it destroys “temporal scarcity.” Temporal scarcity, according to Scheffler, is an implicit assumption of our valuing attitudes. Without it, we would not value basic goods, such as “health, gain, safety, security, and benefit,” because as immortal beings we would have no concept of “loss, illness, injury, harm, risk, or danger” (97). Temporal scarcity, Scheffler continues, is also required for the valuing of more complex goods, such as “creativity, humor, or solidarity” (100). Scheffler claims that immortal human beings could not even be said to make decisions, because human decisions are always made against the background of the limits imposed by time (99).¹²

Scheffler's claim cannot solely concern the conditions under which the experience of valuing or making a decision is possible, because this would reduce the issue to a psychological-state problem. If this were merely a psychological-state problem, then a solution might involve creating the mere *appearance* of temporal scarcity. This would motivate an infinite-life model that incorporates *limited ignorance of the future*,

¹² See also Martha Nussbaum, 1994, 226: “Our finitude, and in particular our mortality ... is a constitutive factor in all valuable things' having for us the value that in fact they have.... The removal of all finitude in general, mortality in particular, would ... bring about the death of value as we know it.”

especially to the extent that such ignorance creates *the seeming possibility of death*. This form of ignorance is, at worst, a minor bad that can be easily outweighed by other positive factors. (Notice that a life in which a person falsely believes that she cannot die can be valuable. Similarly, a life in which a person falsely believes that she can die can be valuable.) With these elements in place, valuing attitudes become possible for eternal beings. Therefore, an argument against the value of infinite life must claim that temporal scarcity is required for our lives to be *of value* — for our valuing attitudes to be “appropriate,” or some related claim.

In light of these considerations, we can see that Scheffler must make a very strong claim if his argument against the value of infinite life is to succeed. He must claim that temporal scarcity is required for a life to be *of value*, and that this temporal scarcity must specifically concern *time alive*. There are good reasons to think that this claim is false.

Let us first take up the worry concerning decision making. The idea is explained in more detail by Richard Wollheim (1984, 265–6). He writes:

“It is natural to think that the point at which deathlessness would most strikingly affect the character of life is at the moment of choice. Many of the reasons that we currently have for choosing this rather than that would be removed.... Questions whether we should do this rather than that would rewrite themselves as questions whether we should do this before or after that, and answers to these new questions would be found by considering not the intensity of our wants or the ends to which they are directed, but the favourable opportunities that the present provides. Weather reports would gain an immense importance in our lives.”

What is being missed here is that eternal beings might be *near biased*; all else equal, they might prefer positive experiences to be near rather than distant, and negative experiences to be distant rather than near. In fact, both Wollheim and Scheffler make the tacit assumption that eternal beings would necessarily be *completely temporally neutral*. Actual people tend not to be temporally neutral, and it is a

mistake to assume that eternal beings would necessarily differ in this regard. There is no necessary connection between the length of one's lifespan and the temporal bias of one's preferences.¹³

By retaining at least a slight near bias infinite beings can be said to make genuine decisions against a temporal framework. Near-biased infinite beings would find it necessary to “consider the intensity of their wants” in so far as they aim to satisfy these wants in the near, rather than distant, future, or in so far as they aim to make trade-offs between lesser but nearer and greater but more distant experiences. In fact, if an eternal being's near bias results in her completely discounting future experiences at some distant point, such as one thousand years in the future, then the temporal framework against which her decisions are made would mirror that of a finite being who expects to live for another one thousand years.

Would the temporal scarcity created by near bias, by itself, also allow eternal beings to retain the valuing attitudes with which Scheffler is concerned? I believe it would. An agent who fully discounts experiences occurring more than one thousand years in the future would possess a temporal framework allowing for concepts like loss, risk, danger, and solidarity. At the very least, if there is some problem with an infinite being possessing these concepts, it would have to be unrelated to the lack of a temporal framework.

There must be some explanation for how valuing attitudes are possible without the prospect of death. Consider the behaviors we observe in children. Niko Kolodny (2013, 168) makes a point about this that is delightful in its combination of insight and obviousness:

“I’m struck by the fact that small children seem to care about many things—such as the attention of their parents, control over their environment, the acquisition of new skills—often quite intensely, even though they don’t have much grasp of mortality, especially not of their own. Of course, one might attribute to them greater implicit sophistication in their beliefs about mortality

¹³ Above we saw that Williams' criticisms of infinite life go awry partly because of his tacit assumption that eternal beings would have greater memory capacities. To assume that eternal beings would be temporally neutral is to make a similar mistake.

or less sophistication in their alleged valuing. But if we take my description at face value, then it suggests that, even if awareness of scarcity, or even temporal scarcity, is necessary for valuing, such awareness induced by awareness of one's own mortality is not.”

There is indeed much evidence that children do not grasp their own mortality.¹⁴ If this is so, then it cannot be true that the temporal scarcity implied by a finite lifespan is required for their valuing attitudes. Whether or not near bias can fully explain the situation is not a question I will take up here, but it suffices to assume that the temporal scarcity created by a feature like near bias, possibly in conjunction with others, does allow for a temporal framework in which valuing attitudes are possible.

A potential response to this argument might start by claiming that near bias is irrational.¹⁵ If near bias is irrational, then Scheffler and Wollheim might be justified in questioning the status of the valuing attitudes that depend on it. The same complaint might be lodged against whatever else these children are doing: perhaps they are just being irrational.

The problem with this response is that there are many potential generators of temporal scarcity, and it is implausible to claim that all are irrational. To find another one we need not look further than the basic alternative. Temporal scarcity in the basic alternative is generated by the presumption that time

¹⁴ See Nichols, 2007.

¹⁵ For reviews of arguments against near bias see Brink, 2010 and Greene and Sullivan, 2015. However, note that traditional arguments against near bias target the rationality of near-biased preferences as they are actually instantiated in human beings, and do not engage with considerations in favor of near bias that might apply to beings with infinite lifespans. On the contrary, it has long been known to researchers of time biases that infinite lifespans can cause problems for temporally-neutral decision makers. One of the most serious problems is what Tjalling Koopmans (1967, 8) called the “paradox of the indefinitely postponed splurge.” The fact that these problems can be avoided by retaining at least a slight near bias is one of the best arguments, I believe, in favor of the rationality of some form of very limited near bias in certain contexts, such as that of infinite life. Such a view would not justify the type of near bias that is observed in actual people.

inside the machine is less valuable than time outside; i.e., what one does while outside the machine “matters more” than what one does while inside. This creates a motivation to accomplish certain things while one is still outside the machine in much the same way that the prospect of death motivates us to accomplish things while we are still alive. If we assume, as argued in Section 3.1, that entering the machine does not destroy the value of our prior activities, then the basic alternative succeeds in removing death while leaving intact the type of temporal scarcity that death creates.

As a final example, consider the possibility of some people living infinite lives while others live finite lives. Those living infinite lives might devote themselves to improving the lot of those living finite lives, and they might rightly take their activities to matter, since i) by hypothesis, finite lives are valuable, and ii) finite lives feature temporal scarcity. Thus, temporal scarcity would be generated through a concern for finite beings. Such a concern would result in infinite beings displaying an active engagement in valuable projects.¹⁶

Are such projects sustainable? Over the course of an infinite life an eternal being might observe an ever-evolving cycle of improvement and deterioration in the average quality of life for finite beings, but each finite life, by hypothesis, remains valuable in itself, as does each instance of improvement of finite life. Must there come a time when there are no more finite beings to aid? No, there is no reason to think that such a time must come. But if it does, an infinite being could always retire to her experience machine.

(4) Conclusion

I have argued that it is possible to avoid the psychological-state and value problems that have featured in the literature on very long life. The upshot of this is the following. First, as a practical matter, we should take seriously the claim that our current lifespans leave room for vast improvement in the case of most individuals (whether longer lives are good all things considered is a different question, which I have not taken up here). Second, as a matter of value theory, we should be skeptical of the claim that death is a

¹⁶ Wolf (1997) argues that this is the key element of a meaningful life.

precondition for a valuable life. It is likely that our valuing attitudes presuppose temporal scarcity and some ignorance of the past and future, but they do not presuppose death.¹⁷

References

- David Blumenfeld, 'Living Life Over Again', *Philosophy and Phenomenological Research* 79 (2009), pp. 357–86.
- Lisa Bortolotti and Yujin Nagasawa, 'Immortality without Boredom', *Ratio* 22 (2009), pp. 261–77.
- Ben Bramble, "On Susan Wolf's 'Good-for-Nothings'", *Ethical Theory and Moral Practice* 18 (2015), pp. 1071–81.
- David O. Brink, 'Prospects for Temporal Neutrality', in C. Callender (ed.), *The Oxford Handbook of Time* (Oxford: Clarendon, 2010), pp. 253–81.
- Mikel Burley, 'Immortality and Boredom: A Response to Wisnewski', *International Journal for Philosophy of Religion* 65 (2008), pp. 77–85.
- William Lane Craig, *Reasonable Faith: Christian Truth and Apologetics*, 3rd edn (Crossway Books, 2008).
- John Martin Fischer, (1994) 'Why Immortality Is Not So Bad', *International Journal of Philosophical Studies* 2 (1994), pp. 257–70.
- John Martin Fischer, 'Free Will, Death, and Immortality: The Role of Narrative', *Philosophical Papers* 34 (2005), pp. 379–403.
- John Martin Fischer, 'Epicureanism about Death and Immortality', *The Journal of Ethics* 10 (2006), pp. 355–81.
- Preston Greene and Meghan Sullivan, 'Against Time Bias', *Ethics* 125 (2015), pp. 947–70.
- Shelly Kagan, 'Immortality', in S. Kagan, *Death* (Yale University Press, 2012), pp. 234–46.
- Niko Kolodny, 'That I Should Die and Others Live', in N. Kolodny (ed.), *Death and the Afterlife* (Oxford University Press, 2013).
- Tjalling C. Koopmans, 'Objectives, Constraints, and Outcomes in Optimal Growth Models', *Econometrica* 35 (1967), pp. 1–15.
- David Lewis, 'Survival and Identity', in A. Rorty (ed.), *Identities of Persons* (University of California Press, 1976), pp. 17–40.
- Thaddeus Metz, 'Recent Work on the Meaning of Life', *Ethics* 112 (2002), pp. 781–814.
- A. W. Moore, 'Williams, Nietzsche, and the Meaninglessness of Immortality', *Mind* 115 (2006), pp. 311–30.

¹⁷ I would like to thank the participants of the 2015 Vancouver Time and Rationality Workshop and the students in Ben Bradley's fall 2015 seminar in ethics and value theory for their constructive discussion of my paper and helpful comments.

- Thomas Nagel, *The View from Nowhere* (Oxford University Press, 1986).
- Shaun Nichols, 'Imagination and Immortality: Thinking of Me', *Synthese* 159 (2007), pp. 215–33.
- Robert Nozick, *Anarchy, State, and Utopia* (Basic Books, 1974).
- Martha C. Nussbaum, 'Mortal Immortals: Lucretius on Death and the Voice of Nature', in M. Nussbaum, *The Therapy of Desire* (Princeton University Press, 1994), pp. 192–238.
- Roy W. Perrett, 'Regarding Immortality', *Religious Studies* 22 (1986), pp. 219–33.
- Samuel Scheffler, *Death and the Afterlife*, ed. N. Kolodny (Oxford University Press, 2013).
- Sydney Shoemaker, 'Personal Identity: A Materialist's Account', in S. Shoemaker and R. Swinburne (eds.), *Personal Identity* (Oxford: Blackwell, 1984).
- Theodore Sider, 'All the World's a Stage', *Australasian Journal of Philosophy* 74 (1996), pp. 433–453.
- Theodore Sider, *Four-Dimensionalism: An Ontology of Persistence and Time* (Oxford University Press, 2001).
- Hunter Steele, 'Could Body-Bound Immortality Be Liveable?' *Mind* 85 (1976), pp. 424–7.
- Richard Taylor. 'The Meaning of Life', in R. Taylor, *Good and Evil* (Prometheus Books, 2000), pp. 319–34.
- Larry S. Temkin, 'Is Living Longer Living Better?' *Journal of Applied Philosophy* 25 (2008), pp. 193–210.
- Leo Tolstoy, *A Confession and Other Religious Writings*, trans. J. Kentish (Penguin Books, 1987).
- Peter Unger, *Identity, Consciousness, and Value* (Oxford University Press, 1990).
- Bernard Williams, 'The Makropoulos Case: Reflections on the Tedium of Immortality', in B. Williams, *Problems of the Self: Philosophical Papers 1956–1972* (Cambridge University Press, 1973), pp. 82–100.
- J. Jeremy Wisniewski, 'Is the Immortal Life Worth Living?' *International Journal for Philosophy of Religion* 58 (2005), pp. 27–36.
- Susan Wolf, 'Happiness and Meaning: Two Aspects of the Good Life', *Social Philosophy and Policy* 14 (1997), pp. 207–25.
- Richard Wollheim, *The Thread of Life* (Cambridge University Press, 1984).