The Ethics of Human Life Extension: The Second Argument From Evolution[1]

CHRIS GYNGELL

Australian National University, Canberra, Australia

One argument that is sometimes made against pursuing radical forms of human life extension is that such interventions will make the species less evolvable, which would be morally undesirable. In this paper I discuss the empirical and evaluative claims of this argument. I argue that radical increases in life expectancy could, in principle, reduce the evolutionary potential of human populations through both biological and cultural mechanisms. I further argue that if life extension did reduce the evolvability of the species this will be undesirable for three reasons: (1) it may increase the species' susceptibility to extinction risks, (2) it may adversely affect institutions and practises that promote well-being, and (3) it may impede moral progress.

Keywords: *Life extension ethics, evolvability, cultural evolution, moral progress*

I. INTRODUCTION

In his paper "The Right and Wrong of Growing Old" Bennett Foddy (2012) identifies two arguments that could be made against human life extension that utilize ideas from evolutionary biology. The first argument from evolution claims that, as the rate of human aging has been optimized by evolutionary forces, we should be sceptical about interventions which alter

[1] A modified version of this article is forthcoming in *The Journal of Medicine and Philosophy*

the aging process. As Foddy demonstrates, this argument is problematic for a number of reasons. First and foremost, most theories regarding the evolution of aging see it as an epiphenomenon a by-product of evolutionary forces and not something that was directly selected. If the rate at which humans age was not selected for by evolution then the claim that the current human lifespan is in some sense optimal loses much of its force.

However there is a second argument against human life extension that draws on evolutionary theory. This argument does not claim that the current human lifespan is in any sense optimal but rather that certain types of life extending interventions may make the species less evolvable, which would be morally undesirable in some way.

Although the second argument from evolution (henceforth SE) is not commonly discussed in the ethical literature, something like it appears to be behind the thinking of many early 20th century scientists and science-fiction writers who wrote about life extension. This is made evident by Adams:

we should note that, remarkably, despite their differences, the diverse works of visionaries, critics, and their successors alike share a consensus about one important point: there is a tension between the interests of the individual and those of the human race as a whole - immortality for the individual may undermine the prospects for the human race. For Wells this is because evolution works best through selective death, the sorting of variants to preserve the best types. Haldane, Huxley and Stapledon grant that we may be able to greatly improve and lengthen human life, but this is

secondary and subordinate to the far more important task of working toward survival and immortality for humankind (Adams 2004, 68).

SE consists of two separate claims: one empirical and one evaluative. The empirical claim is that certain forms of life extension will make our species less evolvable. The evaluative claim is that if our species did become less evolvable, this would be morally undesirable in some way. In this paper I explore both the empirical and evaluative claims of SE.

The empirical claims made by SE are unlikely to hold for interventions which make only minor alterations to lifespan, or which make dramatic changes in only a few individuals. This is because these interventions are unlikely to have the population-level consequences that are needed to affect the evolutionary potential of the entire species. However, interventions which dramatically extend life for large segments of the species could, in principle, have a large enough effect to influence the evolvability of the species. Hence the types of life extending interventions for which SE is most likely to apply are those that make radical changes to *life expectancy* rather than just *lifespan.* [1]

Large increases in life expectancy could reduce the evolvability of the species by slowing the rate of generational turnover – the rate at which new individuals are introduced into our populations. That radical life extension could lead to a reduction in generational turnover has been pointed to by many authors, some of whom are proponents of life extension (Harris 2007). If each individual lives for many times longer than they do now, resource constraints could mean that we need to dramatically reduce the

birth rate as our populations reach their carrying capacity.[2] As Larry Temkin says:

> Even if we grant that the earth could comfortably sustain several times its current population, slowing the aging process would, in a very short time, inevitably carry with it a commitment to slowing the birth rate. This is not a deep philosophical point about the meaning of life or the nature of man. It is a simple practical point…Were we to succeed in stopping our biological clocks, at some point we would only be able to permit new births to offset those deaths due to murders, accident, war and disease (Temkin 2008, 68).

In cases of *radical life extension,* which I define as interventions which increase life expectancy to at least 5000 years and significantly slow generational turnover, the empirical claims made by SE look more plausible. It is these forms of life extension that will be my focus in this paper. In Part II, I look at whether radical life extension would make the species less evolvable through an effect on biological evolution. I argue that radical life extension could, in principle, reduce a population's biological evolutionary potential, but this is unlikely to have a significant effect on the total evolutionary potential of modern human populations. In Part III, I discuss whether radical life extension would make our populations less evolvable through cultural mechanisms. I argue that radical life extension could reduce the capacity of human populations to make adaptive cultural changes, and that this effect may be significant for modern populations. In Part IV, I argue for and qualify the evaluative claims of SE: if radical life extension did reduce our evolutionary potential, this could provide a

reason why such interventions should be viewed as undesirable. In my conclusion, I argue that both the empirical and evaluative claims made by the second argument are plausible, and that SE should be a consideration when assessing the moral desirability of interventions which aim at radical life extension.

II – WILL RADICAL LIFE EXTENSION MAKE THE SPECIES LESS EVOLVABLE THROUGH BIOLOGICAL MECHANISMS?

One way the empirical claims of SE might be true is through an effect on biological evolution. For most species on earth, a significant slowing of the rate of generational turnover would significantly reduce their evolutionary potential. Without a constant influx of new individuals, a population's ability to evolve though natural selection would be significantly curtailed. This is because generational turnover is the most important source of heritable variation for most species of animals. Each new generation introduces novel genes, created by mutation, into the population. Some of these new genes will subsequently result in novel phenotypic traits which can then be subject to natural selection. Similarly, every new generation recombines the existing genetic information contained in a population into novel combinations. Both these processes are important to biological evolution, and are required for most populations to undergo adaptive changes.

Therefore, we would expect any intervention which significantly altered the rate of generational turnover to change the amount of heritable variation that is available to populations of most species of animals. This would subsequently reduce their evolutionary potential. If a population

faces a novel selection pressure, such as changed environmental conditions, a constant supply of new genetic variants can help uncover phenotypes that aid survival in the new conditions. These new variants can then spread through the population, helping the population to adapt to the changed conditions. Hence an intervention which radically extended life expectancy and resulted in individuals being only very rarely replaced could be expected to reduce the evolutionary potential of most populations of animals.[3]

Humans have the same DNA-based inheritance mechanisms as other species of animals, so the effects discussed above are also relevant to human populations. Slowing the rate of generational turnover will lower the rate at which novel genetic variants are introduced into our populations. This may reduce the ability of our populations to make adaptive genetic changes in response to novel challenges. This is one mechanism through which radical life extension may reduce the evolvability of the species, and through which the empirical claims of SE may hold.

One area where this effect may be particularly relevant for modern human populations is with regard to pathogenic threats. This has been suggested by Nick Agar, who states:

negligible senescence is, from an evolutionary standpoint, a bit like asexual reproduction. It gives no opportunity to change the locks on cellular surfaces… (it) may increase our vulnerability to the quadrillions and quadrillions of parasites seeking to infect us (Agar 2010, 124).

Agar is right that dramatically extending lifespan may reduce the ability of populations to biologically evolve in response pathogenic threats. The slower the rate at which novel genetic variants are introduced into a population, the smaller the likelihood that variants which help a population respond to pathogenic threats will be uncovered and the slower they will spread through that population. However it's not clear that this will be a significant problem for modern human populations. This is because we typically respond to pathogenic threats through non-biological mechanisms. As Powell says:

epidemics in human populations can be contained and extinguished in a matter of days, months, or years through cultural–behavioural modifications, including sanitation, vaccination, and the availability of antibiotics; in contrast, the genetic origin and fixation of an immunological adaptation can take hundreds or thousands of years, depending on mutation rates, population structure, selection pressures, and the type of adaptation in question (Powell 2011, 254).

Reducing our capacity to respond to pathogenic threats *biologically* need not increase our total susceptibility to disease, all things considered.[4] This is also true for most other adaptive challenges that modern humans face. The major way our populations adjust to threats in their environments is through changes to phenotype mediated by cultural changes, rather than genetic changes. This has been the primary mechanism through which humans have become so well adapted to their environment across Earth.

Therefore, although radical life extension would in principle reduce the capacity of our species to undergo biological evolution, this does not mean that the empirical claims of SE hold. Radical life extension could reduce the *biological* evolutionary potential of the species, without having a significant effect on the *total* evolutionary potential this species. This is because the major way modern human populations adapt to challenges in their environment is though cultural, rather than biological, mechanisms.

III. WILL RADICAL LIFE EXTENSION MAKE THE SPECIES LESS EVOLVABLE THROUGH CULTURAL MECHANISMS?

If the empirical claims made by SE are to be plausible this will likely be through an effect on cultural, rather than biological, evolvability.
Cultural evolution describes a process through which the distribution of cultural variants[5] (such as ideas, beliefs, and skills), change over time. The propensity of human populations to generate novel cultural variants, and the propensity for adaptive variants to spread and be maintained within populations, helps explain why our species has been so successful. This is made evident by Henrich:

our ecological success, technology, and adaptation to diverse environments is not due to our intelligence. Alone and stripped of our culture, we are hopeless as a species. Cumulative cultural evolution has delivered both our fancy technologies as well as the subtle and unconscious ways that humans have adapted their behavior and thinking to tackle environmental challenges (Henrich 2011).

In this section I want to make some preliminary arguments to show that, in principle, radical life extension could affect cultural evolutionary processes. I will argue that because radical life extension will result in a slowing of the rate of generational turnover, it could influence both a population's propensity to generate novel cultural variants and for adaptive variants to spread through populations.

Generation of Cultural Variants

A regular generational turnover ensures that new individuals are constantly being introduced into a population. This could help generate new cultural information as each individual is raised in a unique cultural environment in which particular ideas and technologies are prominent. As radical life extension will reduce the rate of generational turnover, it could reduce the rate novel cultural variants are generated in human populations. The cultural variants that individuals are first exposed to helps determine how they interpret and respond to information they are subsequently exposed to. This is reflected in the psychological phenomena 'belief dependent realism'. Once individuals have a set of beliefs they become invested in them and this limits their ability to acquire other beliefs (Shermer 2012). This means that the order in which we acquire cultural variants is important. When we acquire certain ideas early in development we start to see the world in certain ways. It affects how we search for and interpret information, and influences which cultural variants we will subsequently acquire.

A constant generational turnover ensures that naive minds are constantly developing in unique cultural environments and are exposed to different cultural variants early in development. This means that, over time, populations have a large number of different individuals exist in them, each of whom are disposed to develop different cultural variants. Therefore, as radical life extension will slow the rate of generational turnover, it may reduce the variance in ideas and skills that are available to populations over time.

Slowing the rate of generational turnover may also reduce the rate at which cultural variants are transmitted between individuals. This could slow the rate at which novel cultural variants are generated. Some cultural variants only need to be acquired once, after which individuals do not need to re-acquire them. Therefore, in a population where individuals are rarely or never replaced, some cultural variants will not need to be transmitted between individuals often. For example, once everyone in a population who wants to know how to fish, can fish, there is no need for fishing skills to be transmitted between individuals. In contrast, in populations with a regular generational turnover, naive unskilled individuals are constantly introduced into the population. This means skills and ideas need to be continually re-transmitted between individuals.

The process is likely to generate novel cultural variants for the simple reason that the transmission of cultural variants from one individual to another is "inherently error-ridden" (Aunger 2010, 3). Individuals have only imperfect access to the cultural information that is stored in the brains of others and have to infer what ideas or beliefs others have. Because we don't have direct access to the representations of others, the process of

transmitting cultural variants from one individual to another can be messy – errors arise and this can produce novel cultural variants.

The tendency for increased transmission of cultural variants to result in the increased generation of novel variants can be seen in the case of skills. Often individuals will learn skills through a combination of social mechanisms and individual trial and error learning – a process known as 'apprentice learning' (Sterelny 2012). This process is likely to lead to variations in the skills realised in the learner, due to differences in their characteristics and their learning environment. For example, an individual trying to learn how to perform a tennis serve may end up with a slightly different serve to their teacher. This is because when they practice and engage in trial and error learning; they may find that slight variants of the serve work best for them, due to differences between themselves and their teacher or differences in their learning environment. Therefore the greater number of individuals with unique characteristics coming through a population and re-learning skills in unique learning environments, the more variations in those skills we can expect. These new variants can affect the trajectory of cultural inheritance in the next generation, because successful variants will act as models for subsequent generations.

Spread of Cultural Variants

Another aspect of cultural evolution that could be affected by radical life extension is the propensity of cultural variants to spread through a population. A regular generational turnover provides a steady stream of children into a population, and children may be particularly well suited to

copying cultural variants from others. One reason why the presence of children in a population may help cultural variants spread is that the brains of children are particularly malleable. As Jablonka and Lamb (2006, 166) state "there seems to be a special 'window of learning' for some types of behaviour – a window that is wide open early in life and gradually closes as the individual matures". This is reflected in the idea of critical, or sensitive, periods in development. The neural substrates required to perform certain skills do not develop normally if appropriate stimulation is not received during a restricted time period in development. For example – musicians who began practising before the age of seven have been shown to consistently outperform older trained musicians, even when factors like total years trained are controlled for (Watanabe, Savion-Lemieux, and Penhune 2007). These differences in ability are mirrored by differences in the brains of the these musicians, with those who began practising before the age of seven showing much greater levels of myelination in areas of the brain associated with musical ability. Other critical periods have been postulated for visual perception (Lewis and Maurer 2005), language acquisition (Lenneberg 1967), and auditory processing (Sharma and Campbell 2011).

Older individuals may therefore have difficulties in learning complex skills they observe in others, because their brains have passed various 'critical' periods, and are not as flexible as the brains of children. If an individual develops a complex skill in a population consisting solely of mature individuals it may be unlikely for it to spread as it will be difficult for other members to adopt it.[6]

Besides being less *able* to adopt novel cultural variants, older individuals may also be less *motivated* to adopt them. One reason that adults may be less motivated than children to adopt new ideas and beliefs was touched on above – once individuals acquire certain ideas and beliefs they become invested in them and have trouble giving them up. This is reflected in the idea of 'confirmation bias' – a tendency to seek and interpret information in a way that confirms one's beliefs (Nickerson 1998). For example, when reading about political issues like gun control and affirmative action, individuals have been shown to prefer sources that affirm their existing attitudes and to interpret ambiguous evidence as supporting their existing position (Taber and Lodge 2006). Adults are sometimes unmotivated to change their beliefs, because they do not perceive their current beliefs as deficient in any way. As children tend to not have as many fully formed beliefs as adults, they are not as prone to confirmation bias. As a result they may be more likely to acquire novel ideas and beliefs. A steady stream of children may therefore help novel ideas spread through populations.

The above is just a brief overview of some of the ways in which radical life extension could affect cultural evolutionary processes. By reducing the rate of generational turnover, radical life extension will reduce the number of individuals who are introduced into a population. This may reduce variation in the cultural information that is available to a population, as each individual has unique characteristics and a unique cultural background and hence will have their own propensity to develop cultural variants. Further, reducing the rate of generational turnover also reduces the number of children who live in a population. This may affect the

propensity of cultural variants to spread in a population as children have characteristics which make them particularly disposed and able to copy cultural variants.

If radical life extension were to reduce the number of cultural variants available to a population, and also reduce the propensity of adaptive cultural variants to spread, then it is likely it would reduce the cultural evolvability of that population. Culture may be less likely to change in these populations, as they will consist nearly solely of mature individuals who are invested in their cultural variants and who are less likely to adopt novel cultural variants. Hence radical life extension could have a stagnating effect on culture. This provides one mechanism through which radical life extension could make populations less evolvable.

Therefore, although the specific cultural evolutionary consequences of radical life extension will be highly dependent on the specific details of the intervention and the population it occurs in, it is plausible that in certain conditions radical life extension will reduce cultural evolvability. The second argument from evolution, therefore, cannot be dismissed on purely empirical grounds. In cases of radical life extension, there are plausible ways these interventions will make us less evolvable. Whether this provides any reason to think that radical life extension will be morally undesirable is the focus of the next section.

IV ETHICAL IMPLICATIONS

In this section I will examine the evaluative claims of SE. If an intervention did make a population less evolvable, would this be morally undesirable? I

will look at three reasons why interventions which reduce evolvability could be undesirable. Reduced evolvability could (1) increase the species' susceptibility to extinction risks, (2) reduce well-being for individuals, and (3) impede moral progress. These factors point to possible costs associated with interventions which would result in radical life extension.

Increased Susceptibility to Extinction Risks

One reason a reduction in evolvability may be undesirable is that it will increase our susceptibility to extinction risks and make the species less likely to survive well into the future. If radical life extension adversely affects our ability to generate novel ideas and skills, we may be less able to adapt to environmental and societal challenges. This may make the species more likely to become extinct as a result.

If radical life extension did increase the species' susceptibility to extinction risks, this could be morally undesirable. Some theorists claim that we have either instrumental or intrinsic reasons to value the continued existence of the species. For example (Bostrom 2002, 4) describes reducing our susceptibility to extinction risks as a "global public good". Sidgwick (1907, 487) describes the extinction of mankind as the "greatest conceivable crime", a view with which Parfit (1987) sympathises. Agar (2010) defends a theory of goods called 'species relativism' which also implies we have reasons to try to prevent species extinction.

In this paper I will not provide a thorough investigation of whether we actually have good moral reasons to value the continued existence of the species. I will merely note that some theories do imply that we have such

reasons.[7] Therefore, if radical life extension did reduce the evolvability of the species, then under these theories of goods, it would provide one reason why it would be undesirable.

Effects on well-being

Another reason that radical life extension could be undesirable is that it will reduce well-being for individuals. Specifically, a reduction in the rate of generational turnover – which can lead to reduced cultural evolvability – may adversely affect various practises and intuitions which promote well-being. Two broad areas which contribute to well-being and may be affected by generational turnover are the Arts and Sciences.

It's plausible that the Arts (music, literature, and so forth) contribute to well-being for individuals living in a population. If these disciplines were adversely affected by interventions which reduced generational turnover, this would point to ways in which such interventions may reduce well-being.

Intuitively, generational turnover seems important to the Arts. Each new generation brings fresh perspectives and new ideas to various artistic disciplines, and this seems to contribute to the diversity in artistic styles and the richness we observe in these disciplines.

For example, music in populations with a high rate of generational turnover may be richer than in populations with a slow rate of generational turnover. This is because populations with a high rate of generational turnover will, over time, have more musicians than populations with a

slow rate of generational turnover. This is simply a result of the fact that populations with high generational turnover will have many people cycling through them- and by extension more people who choose to become musicians. Similarly a regular generational turnover ensures that individuals who choose to take up music are constantly growing up in a unique cultural environment, which influences the music they ultimately create. These factors mean we would expect more variety in music in populations with a high level of generational turnover than in populations that have a very slow rate of generational turnover.

Of course, there could be some ways in which radical life extension may benefit music. For example it's possible that in a world where musicians live for thousands of years, they will become much better musicians and be able to develop richer and more complex music as a result. This is a possible benefit, but it's not clear that we expect musicians to continually get better and better as they age. Lots of experience does not seem like a pre-requisite for creating good music and many great musicians have died young. Still, there are possible advantages to slowing generational turnover if this allows musicians to acquire extra experience. These would need to be balanced and weighed against the costs associated with a reduced number of musicians in total, and the likely reduction in diversity this will cause. It's possible that at some point, the gains associated with letting specific musicians live longer would not offset the costs of having fewer musicians in total over time.

The same can be said for other artistic disciplines. More artists in total over time, and more artists growing up in unique cultural niches, may benefit many artistic fields. Hence a risk of interventions which result in a

reduction in the rate of generational turnover could be that they have an adverse effect on various artistic disciplines. This also has the potential to negatively affect well-being, as having a diverse and vibrant Arts culture can be seen as contributing to the well-being of individuals.

Another area which may be adversely affected by reduced generational turnover is the sciences. It seems plausible that progress in science is aided by a regular generational turnover which ensures a continual supply of new scientists. New scientists can bring different perspectives to problems and this helps drive science forward. Weisberg and Muldoon (Forthcoming) suggest that one factor which helps promote scientific progress is a mix of 'maverick' strategies, which pursue novel research agendas, and 'follower' strategies which pursue tried and proven research agendas. A mix of these strategies is the most efficient way to divide cognitive labour when searching through epistemic landscapes. It's possible that new and younger scientists may be more likely to align with maverick strategies. Often younger scientists are perceived as the ones more likely to make fundamental leaps in the field, as made clear by Kuhn:

almost always the men who achieve these fundamental inventions of a new paradigm have been either very young or very new to the field whose paradigm they change (Kuhn 1962, 89-90).

Generational turnover may help supply scientists who pursue novel research agendas, in part because it supplies scientists who have been raised in unique cultural niches as discussed above. This may make these

scientists more likely to stumble upon previously undiscovered peaks in the epistemic landscape and help drive science forward.
Similarly, a constant supply of new scientists without preconceived ideas and beliefs about their field may also help scientific progresses by helping good ideas spread and gain prominence. If novel ideas arise which challenge conventional scientific wisdom, they may not spread in a population consisting solely of mature scientists. This is because mature scientists are more likely to suffer from the types of conformational biases discussed above. Having preconceived beliefs about a scientific field may interfere with one's ability to fairly assess ideas that are contrary to accepted wisdom. This has been indicated by the German physicist Max Planck who once observed "a new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it" (Wolinsky 2008, 1-2).

Of course it is possible that radical life extension may also have beneficial effects on science. For example, it's possible that scientists who live for thousands of years may be able to continually learn from their experiences, and generate more complex and better ideas about science as a result. While this is a possible benefit of radical life extension, to a certain degree, this occurs anyway. Present day scientists learn from the experiences of previous generations of scientists, and can generate more complex ideas as a result. It's unclear what benefits would accrue from letting *single* individuals acquire experiences for thousands of years which could not be achieved through intergenerational collaboration.

One possible consequence, then, of reducing cultural evolvability by slowing generational turnover is the slowing of scientific progress. This could reduce well-being for individuals living in a population as scientific discoveries often contribute to well-being. Advances in medicine, for example, are thought to make people's lives better by improving health. Such advances may be fewer and further between in populations with slower generational turnover.

One possible way that interventions which reduce cultural evolvability may be undesirable is by having an adverse effect on institutions and practices that promote well-being. Two broad areas which contribute to well-being and may be adversely affected by reduced cultural evolvability are the Arts and Sciences. A regular generational turnover, providing novel cultural variants, helps ensure these disciplines flourish.

Of course this does not necessarily mean that, *all things considered*, well-being will be reduced in populations with a reduced generational turnover. When looking at interventions aimed at radical life extension specifically, people may prefer to live in populations with slower scientific progress and less diverse artistic disciplines, if it means they can live for thousands of years. Whether, and to what extent, a radically extended lifespan will improve individual well-being has been looked at in detail by others (Williams 1981; Temkin 2008). In this section I merely want to note that there are some ways in which being less evolvable could negatively impact on well-being, and these costs would need to be weighed against any well-being benefits associated with having a radically extended life-span.

Moral Progress

A final reason that reduced cultural evolvability may be undesirable is that it may impede moral progress. Moral beliefs are influenced by confirmation bias and once formed they are difficult to change (Haidt 2001). This means that something very close to Planks' quote about progress in science above could apply to moral progress. That is 'new moral ideas do not triumph by convincing their opponents and making them see the light, but rather because their opponents eventually die, and a new generation grows up that is familiar with them.'

Peter Singer (2011) describes moral progress as 'expanding circles of moral concern'. While early humans were only concerned with the well-being of their immediate kin and group members, society has been slowly evolving so that individuals have begun to care for larger and larger portions of humanity. This was initially limited to concern for individuals in neighbouring groups, then to individuals in larger national and ethnic groups, and is in the process of expanding to all of humanity and non-human animals. This type of progress has been epitomised over the last century with the increasing recognition of rights for disadvantaged groups. This began with the women's rights movement and the African-American civil rights movement, and can be seen more recently in the lesbian, gay, bisexual, and transgender (LGBT) social movement and the animal rights movement.

Generational change can be seen as playing an important role in these types of rights movements. The women's rights movement is strongly associated with the 'roaring twenties', a period of significant social and moral change that occurred during the 1920's. Today, younger generations are much

more supportive of LGBT rights, and show greater concern for animal welfare, than older generations (Smith and Son 2013; Wilke and Saad 2013). Presumably this is partly due the fact that once individuals have fully formed beliefs about whether a particular group is deserving of equal rights, it may be difficult for them to change this belief, due to confirmation bias. A regular generational turnover can therefore be seen as important to moral progress as it provides a continuous supply of individuals without preformed moral beliefs.

Therefore, if we were to slow the rate of generational turnover, this may reduce the opportunities our populations have for moral reform. Our populations will consist nearly entirely of individuals with fully formed moral beliefs who have difficulty changing these beliefs.

In this section I have outlined three ways in which interventions which reduce evolvability may be undesirable. Hence I have tried to show that the evaluative claims made by SE are plausible. If the species as a whole becomes less evolvable then this may be undesirable as it will increase our susceptibility to extinction risks, negatively impact some institutions and practices that improve well-being and impede moral progress.

However it needs to be stressed that this does not mean that, all things considered, radical life extension, or any intervention which did reduce evolvability, would be morally undesirable. Having a vastly extended lifespan may be highly desirable for a host of other reasons that outweigh concerns related to reduced evolvability. Still, the costs relating to reduced evolvability need to be considered when assessing the moral desirability of interventions which aim at radical life extension.

V. CONCLUSION

In his paper 'Is Living Longer Living Better?', Larry Temkin asks:

> would an outcome in which society's members lived forever be ideal? Would it be as good as an outcome where society's members were continually replaced every 100 years, or perhaps every couple of centuries? (Temkin 2008, 193)

In this paper I have not tried to answer Temkin's question, but rather attempted to show that a specific argument is applicable to it. The second argument from evolution claims that interventions which radically extend lifespan and slow generational turnover will make our species less evolvable, and that this will be undesirable in some way. I have argued both these claims are plausible. A regular generational turnover provides populations with a constant source of genetic and cultural novelty which helps with biological and cultural evolvability. If our populations were to become less evolvable, there are several ways in which this may be undesirable – it may reduce well-being for individuals, impede moral progress and make us more susceptible to extinction risks. Hence the second argument from evolution provides some reason why we would prefer the second alternative in Temkin's questions – an outcome in which society's members were continually replaced every 100 years, or every couple of centuries – over one in which society's members live indefinitely.

NOTES

[1] In this paper, I follow (Post and Binstock 2004) and use the term lifespan to refer to the maximum amount of time that a member of a species has lived, and life expectancy to refer to the average amount of time that a member of a species lives.
[2] Aside from these practical considerations, there may also be psychological reasons why radical life extension would result in a lowering of the birth rate - eternal youth may reduce the desire of individuals to have children (President's Council on Bioethics 2011).
[3] This idea forms the basis of group-selection or "evolvability" theories of aging (Longo, Mitteldorf, and Skulachev 2005; Goldsmith 2008). The claims these theories make regarding the evolutionary origin of aging are highly controversial, mainly because they appear incompatible with natural selection (De Grey 2007). However even if aging did not specifically evolve as a way to help populations adapt, it could still have this effect. If aging helps keep the rate of generational turnover high, which in turn introduces novel genetic variants into a population, this will help populations adapt to changing conditions.
[4] It should be noted that in the future genetic engineering technologies may mean we can alter the components of our innate immune system directly. This would mean that populations could respond to disease threats biologically without the need to wait for mutations created by generational turnover.
[5] In this paper, I will follow Richerson and Boyd (2008) and use the term cultural variants to refer to the units of cultural information *i.e.* "ideas", "beliefs", "values", and "skills". However I remain neutral as to the precise mechanisms of cultural evolution. By using their terminology I do not wish to endorse Richerson's and Boyd's 'dual inheritance' model.
[6] It possible that by the time radical life extension is technologically feasible, it may also be feasible for individuals to re-engineer their brains so that they can effectively learn new skills in spite of the fact that they have passed critical development periods. If this is the case, then radical life extension may have little effect on population's ability to acquire new skills.
[7] See Gyngell (2012) for a more in-depth discussion.

REFERENCES

Adams, M. 2004. The Quest for Immortality: Visions and Presentiments in Science and Literature. In *The Fountain of Youth*, edited by S. Post and R. Binstock. Oxford: Oxford University Press.

Agar, N. 2010. *Humanity's End: Why We Should Reject Radical Enhancement* Cambridge, Massachusetts MIT Press.

Aunger, Robert. 2010. Three Roads to Cultural Recurrence. In *Evolutionary Psychology and Information Systems Research*, edited by N. Kock. New York: Springer.

Bioethics, President's Council on. 2011. *Beyond Therapy: Biotechnology and the Pursuit of Happiness*: General Books.

Bostrom, N. 2002. Existential risks: analyzing human extinction scenarios and related hazards. *Journal of Evolution and Technology* 9 (1):1-30.

De Grey, A. D. 2007. Calorie restriction, post-reproductive life span, and programmed aging: a plea for rigor. *Annals of the New York Academy of Sciences* 1119:296-305.

Foddy, Bennett. 2012. The Right and Wrong of Growing Old: Assessing the Argument from Evolution. *Philosophy & Technology* 25 (4):547-560.

Goldsmith, T. C. 2008. Aging, evolvability, and the individual benefit requirement; medical implications of aging theory controversies. *Journal of Theoretical Biology* 252 (4):764-8.

Gyngell, Chris. 2012. Enhancing the Species: Genetic Engineering Technologies and Human Persistence. *Philosophy & Technology* 25 (4):495-512.

Haidt, Jonathan. 2001. The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* 108 (4):814-834.

Harris, John. 2007. *Enhancing evolution : the ethical case for making better people*. Princeton, NJ: Princeton University Press.

Henrich, J. 2011. A cultural species: How culture drove human evolution. *Psychological Science Agenda* [On-line]. Available: http://www.apa.org/science/about/psa/2011/11/index.aspx (accessed October 9, 2013)

Jablonka, E. and M. Lamb. 2006. *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge, Massachusetts: MIT Press.

Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press

Lenneberg, E.H. 1967. *Biological foundations of language*. Oxford: Wiley.

Lewis, T. and D. Maurer. 2005. Multiple sensitive periods in human visual development: evidence from visually deprived children. *Developmental Psychobiology* 46 (3):163-83.

Longo, V. D., J. Mitteldorf, and V. P. Skulachev. 2005. Programmed and altruistic ageing. *Nature Reviews Genetics* 6 (11):866-72.

Nickerson, R. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2 (2):175-220.

Parfit, D. 1987. *Reasons and persons.* Oxford: Oxford University Press.

Post, S. and R. Binstock. 2004. Introduction. In *The Fountain of Youth*, edited by S. Post and R. Binstock. Oxford: Oxford University Press.

Powell, R. 2011. The Future of Human Evolution. *The British Journal for the Philosophy of Science* 63 (1):145-175.

Richerson, P.J. and R. Boyd. 2008. *Not By Genes Alone: How Culture Transformed Human Evolution*. Chicago: University of Chicago Press.

Sharma, A. and J. Campbell. 2011. A sensitive period for cochlear implantation in deaf children. *Journal of Maternal-Fetal and Neonatal Medicine* 24 Suppl 1:151-3.

Shermer, Michael. 2012. *The Believing Brain: From Ghosts and Gods to Politics and Conspiracies---How We Construct Beliefs and Reinforce Them as Truths*. New York: St. Martin's Griffin.

Sidgwick, H. 1907. *The methods of ethics*. Seventh ed. London: Macmillan.

Singer, P. 2011. *The Expanding Circle: Ethics, Evolution, and Moral Progress*. Princeton, NJ: Princeton University Press.

Smith, T. and J. Son. *Trends in Public Attitudes about Sexual Morality* [On-line] Available: http://www.norc.org/PDFs/Trend_Report_Sexuality.pdf (accessed October 9, 2013)

Sterelny, K. 2012. *The Evolved Apprentice: How Evolution Made Humans Unique* Cambridge, Massachusetts: MIT Press.

Taber, Charles S. and Milton Lodge. 2006. Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science* 50 (3):755-769.

Temkin, Larry S. 2008. Is Living Longer Living Better? *Journal of Applied Philosophy* 25 (3):193-210.
Watanabe, D., T. Savion-Lemieux, and V. B. Penhune. 2007. The effect of early musical training on adult motor performance: evidence for a sensitive period in motor learning. *Exp Brain Res* 176 (2):332-40.
Weisberg, M. and R. Muldoon. Forthcoming. Epistemic Landscapes and the Division of Cognitive Labor. *Philosophy of Science*.
Wilke, J. and L. Saad. 2013. *Older Americans' Moral Attitudes Changing* [On-line]. Available: http://www.gallup.com/poll/162881/older-americans-moral-attitudes-changing.aspx (Accessed 9 October 2013)
Williams, B. 1981. *Moral Luck*. Cambridge: Cambridge University Press.
Wolinsky, H. 2008. Paths to acceptance. The advancement of scientific knowledge is an uphill struggle against 'accepted wisdom'. *EMBO reports* 9 (5):416-8.