HOW LONG OUGHT WE TO LIVE? THE ETHICS OF LIFE EXTENSION

By

ADRIAN BUNN

(BA hons.)

A thesis submitted to Charles Sturt University in fulfilment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Charles Sturt University

(March 2014)

CONTENTS

Certificate of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the award of any other degree or diploma at Charles Sturt University or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by colleagues with whom I have worked at Charles Sturt University or elsewhere during my candidature is fully acknowledged.

I agree that this thesis be accessible for the purpose of study and research in accordance with the normal conditions established by the Executive Director, Library Services or nominee, for the care, loan and reproduction of theses.

Adrian Bunn

ACKNOWLEDGEMENTS

This thesis would not have been possible without the friendship, support, and patience of my principal supervisor, Dr. Daniel Cohen, whose steadfast encouragement and enthusiastic conversations have contributed significantly to the content and quality of the thesis. Similarly, I would like to express my gratitude to the help and support of my former principal supervisor, Dr. Steve Matthews, who set me on the right path. The advice given by my secondary supervisor, Prof. John Weckert, has also been invaluable to the completion of this thesis.

I would also like to acknowledge the supportive and stimulating environment of Charles Sturt University and its philosophy faculty, as well as the special research centre, The Centre for Applied Philosophy and Public Ethics and its members at ANU, where I spent my first years as a new PhD candidate. Finally, I would like to thank my parents for their encouragement and financial support throughout my PhD candidature, and for their tolerance of their son's extended time in academic study, first as an undergraduate and then as a postgraduate student.

ABSTRACT

In the future we will have technology capable of extending human lifespans beyond their current maximums, making the issue of longer lives a pressing one for the species. Some have claimed that humans living longer lives is not self-evidently a good thing. Others have argued that developing technologies to radically extend the normal human lifespan would be wrong because the negative consequences for societies and for future generations would outweigh the benefits of having the technology. I argue in this thesis that developing life extending technology would be permissible.

I proceed first by examining two prominent sets of arguments from Leon Kass (2004) and Bernard Williams (1973) that the personal desire for a profoundly longer life is an irrational one. Their arguments attempt to demonstrate that living too long would be antithetical to a meaningful and attractive existence for an individual. I closely analyse both Kass's and Williams's arguments and ultimately find that they do not demonstrate that the desire for profound personal survival is irrational or that survival is necessarily a bad thing for the individual.

Building on this conclusion, I argue that life extension might be a good thing by showing why death can be a harm to the one who dies and why a longer life is better than a shorter one, other things being equal. I do this first by defending and modifying views about why death can be bad for the one who dies, such as Nagel's 'deprivation account' and Williams's account of the importance of 'categorical desires'. I also defend a 'non-episodic' view of personal experience and desire satisfaction against the Epicurean 'episodic' view of personal time to reveal why a longer life, all else equal, is better than a shorter one.

If survival is not necessarily a bad thing for the individual and a premature death can be a harm, then life extension—as far as the individual is concerned—must be a good thing. I assume these

conclusions as I proceed to argue that developing life extending technology would be permissible. Focusing exclusively on an argument from Peter Singer (1991) that developing a life extension drug that doubled the average human lifespan to 150 years would be impermissible—because doing so would bring about a future world with lower total and average wellbeing—I highlight problems with the moral principles that guide Singer to his conclusion and propose a more acceptable version of average utilitarianism, which, if applied to the case of developing a life extension drug, demonstrates that developing a drug that doubled the human lifespan would result in an outcome with higher average happiness per person, making the decision to develop life extension permissible on a consequentialist view.

Proceeding forward on the assumption that developing life extending technology would be permissible, I challenge the argument that the only just distribution of life extending technology is to provide everyone equal access. I present the case that life extending technology would be most fairly distributed under a certain interpretation of the Rawlsian difference principle of resource distribution in a way that gave priority to those least advantaged with regard to life expectancy. My defence of this prioritarian version of the difference principle is also compatible with the view that older persons can sometimes deserve priority for life extending interventions over younger persons, thus challenging the 'fair innings' intuition that we ought to, in principle, give priority for life extension interventions to younger persons over older persons when in competition for these resources.

The conclusions of this thesis reveal some misconceptions about the value of longer lives for individuals and of developing life extending technology. What we need is more detailed philosophical analysis of the issues relating to humans beings living longer lifespans and the consequences of having technology capable of radically extending human lives than has been devoted to these issues up to this point.

CHAPTER 1

INTRODUCTION

Why talk about life extension? In the future we will have technology capable of prolonging human lives beyond their current maximums. Perhaps not in the immediate future, as I will suggest in Chapter 2, but soon enough to make the issue a pressing one for the species. The possibility of dramatically extending human lifespans inspires in some a sense of wonder, in others anxiety, but otherwise invites us to answer some basic questions about whether human beings living substantially longer is self-evidently a good thing and whether developing life extending technology would be permissible. This thesis answers the second question by developing arguments to support the claim that life extending technology is permissible.

What quickly becomes apparent is that there are many revealing and interesting questions surrounding life extension that are not strictly normative in nature. For example, an individual might ask whether the desire for a profoundly longer life is a rational one, that is, might living too long involve suffering or be self-defeating in some other way? This question I address in Chapter 3. Alternatively, one might ask whether death is a bad thing for the one who dies, such that avoiding death by extending one's life is a worthwhile pursuit, maybe even rationally obligatory. I examine this question in Chapter 4. We human beings might also collectively ask ourselves whether individuals existing for longer would in itself be a good thing, and whether it would be the wisest course of action to develop life extending interventions.

The central goal of this thesis, however, will be to answer a normative question: Should we—societies or governments—develop life extension? This is a question that any ethical analysis of life extension needs to answer. I will start to give my answer to this question in Chapter 5, but along the way I will address many of the other questions surrounding life extension, and it will become evident how some of these questions collapse into each other.

1.1 What is our aim in extending human lives?

What do we mean when we say we want to extend the human lifespan? The majority of the empirical literature about life extension suggests the typical aim is to intervene in human ageing: to cure humans of old age. This will likely be the central path to humans experiencing longer lives in the future. What ageing researchers, called biogerontologists, mean by 'curing human ageing' is to make the human organism exempt from experiencing the biological ageing process that occurs after the organism has reached maturation, resulting in the decline of an organism's capacity for cellular replication and the deterioration of the health and functioning of the organism leading to the organism's eventual death. This, of course, is only one potential end of the human organism, and death by 'old age' only occurs in the event that some external harm or other doesn't cause death first.

However, the scientific enterprise to cure humans of old age is tangled up in controversial empirical and conceptual notions. Let's start with a simple one: the ambiguity in the terms 'ageing' and 'old age'. By 'ageing' we could mean either (a) the passing of time, or (b) a biological process. Likewise with 'old age' we could mean, (a) time has passed, or (b) a biological process is occurring or has completed. What biogerontologists mean is (b) in both instances. They want to intervene in a biological process, called ageing, of which old age is the result of this process or sometimes just the state at the end of this process, something we ideally want to avoid, or at least postpone for as long as possible.

A further ambiguity with the term ageing, whose disambiguation is more important than the last, is as follows: in the context of curing ageing, anti-ageing researchers aim to cure 'adult ageing', or the deterioration that is the result of changes post-maturation. Anti-ageing researchers do not want to prevent nascent humans from developing or changing into adult humans. This is not the harmful or 'bad' kind of ageing but rather a process of useful developmental changes. We don't want to create long-lived 8 year-olds! So there seems to be a distinction between beneficial ageing and harmful ageing. The target of anti-ageing science is the ageing process that ceases to result in beneficial change and becomes harmful deterioration after the organism reaches its state of maturation. Thus I will assume in the thesis that curing ageing means curing the detrimental biological process of adult human ageing. In Chapter 2 I will

review the dominant views about why humans experience this sort of ageing at all and what this means in terms of its possible retardation or reversal.

The aim of curing ageing also presupposes that ageing is a pathology and not a normal process. However, this portrayal of human ageing is notoriously controversial in the empirical literature (see, for example, Caplan 1981; 2005; de Grey *et al* 2002a; Engelhardt 1979; Hayflick 1998). Nonetheless, anti-ageing medicine is founded on the idea that ageing is a disease and that what we die from is old age, a result of that disease, assuming some kind of external harm doesn't get us first. If we want to extend lives, then we need to prevent humans from experiencing the non-intrinsic, detrimental process that is adult human ageing. But is ageing the disease that people die from? Do people, strictly speaking, die from old age, or do they die from anything else but old age, for example, other diseases like cancer or heart disease? Yet another interpretation of old age is to simply think of it as a threshold or a limit, one that nobody actually reaches because they die from something else before they can. If we agree that what people die from is old age, then that means old age is a disease and what we aim to cure is old age. Or, if nobody ever dies from old age, understood as a kind of limitation, then our aim must be either to cure other diseases that cause people to die or to address our vulnerability to the diseases that cause death. It seems that either course of action takes us down the road of ultimately trying to remove causes of death for humans, which, in essence, is life extension.

If we go the first route to life extension—curing old age—our aim is to cure just one cause of death, that is, death from old age. If we go the second route—curing other diseases—our aim is to cure all other diseases one by one, which would take longer and would also significantly increase our chances of survival. However, even after we cured all other diseases (an enormous achievement in itself!) we would still die from accidents or other external harm. Thus whichever route we take to life extension would also need to prevent humans from dying from external causes of death as well for interventions to be truly effective in substantially prolonging lives.

So even if ageing is a distinct disease in its own right, we would still need to cure particular human diseases one at a time. The second path to life extension requires us only to cure the other diseases. But we might also aim to make people less vulnerable to other diseases because the older people get the more susceptible they are to disease. So a third way might be not to cure ageing as such but rather to make old people less susceptible to disease. Essentially we want to do here is make old people as strong and

healthy as young people, managing disease rather than curing it. On both second and third paths to life extension, people will still age, in the sense of continuing to live and accumulating their number of years alive, but they will not experience physical or mental decline.

The third approach of managing disease rather than curing it is directly related to the dispute among biogerontologists about whether their goal should be the compression of morbidity, i.e. compressing the time one spends frail at the end of life, or directly intervening into the ageing process (See: de Grey 2007; Juengst 2004). Some, biogerontologists, such as de Grey, believe that compression of morbidity is the second best approach, and that it stems from traditionalist views that we should not intervene directly into ageing but our only requirement is to compress morbidity. de Grey seems to be right, that if biogerontology's ultimate goal is to remove human suffering associated with ageing, then directly tackling the ageing process will yield greater results (if ageing is a single process and can indeed be manipulated: see section 2.1) than merely managing old age disease and debility.

Part of the disagreement among biogerontologists' on this issue is also about avoiding the dispute over the moral praiseworthiness of 'enhancement' as opposed to 'treatment'. The compression of morbidity can easily be viewed as treatment of a normal end of life process, while life extension is commonly felt to be enhancement of normal functioning. Yet those like Bostrom and de Grey argue that the goal of compression of morbidity is harmfully diverting energies away from the more beneficial goal of biogerontology of intervening in ageing directly and thereby 'saving lives' in virtue of denying ageing, the biggest taker of lives, its victim (Bostrom 2005a; de Grey 2007).

However, one might ask whether it matters what scientists aim to cure just so long as the result of their intervention is the postponement or complete avoidance of causes of death for the human organism. One might take the perspective that our ultimate aim is to reduce human suffering and increase human happiness, which will also be the result, at least for human individuals, of curing either ageing or other causes of death and decline. Thus the simple answer to the question of whether we should cure ageing and other related diseases is "yes" because doing so would be beneficial to humans and not doing so would be harmful.[1] But there remain complicated arguments

[1] Again, see articles by Bostrom and de Grey.

that we should not remove causes of death for the human organism, ageing especially, which I will introduce later in the Introduction.

Clearly it would be of little or no value to anybody, most of all the individual, if the aim were simply to extend the human lifespan with no removal of decline or debility. This is not the aim of those who promote the extension of human lives through curing human ageing. I will take for granted throughout the thesis, then, that the goal of curing ageing, and of other biological life extending strategies, is the extension of the health and functioning of the human organism together with the prolongation of the lifespan, or what has been called the extension of the 'healthspan' (de Grey 2005; Farrelly 2010; Gems 2003).

1.1.1 Life extension and living forever

What if we cured ageing or some other life extending strategy were successful? Does this mean we would live indefinitely? Curing ageing, or the continual application of life extending interventions, while not likely to transmute human existence into athanasia (a 'state without death'), might allow us to say we were 'contingently' immortal, that is, ageless but still capable of dying from external harm. So long as the ultimate goal of removing all causes of death for humans (including vulnerability to all external damage) remains incomplete so does immortality remain an unattainable condition for human beings.[2] The only way humans could possibly achieve this is through eliminating all causes of death for human beings, but that such a goal seems unattainable. To appreciate why the notion of 'immortality' is so problematic when applying it to a human life (or any sentient life) we need to look a little closer at the notion of immortality.

On a philosophical level, the idea of necessary immortality is conceptually problematic.[3] If immortality is, strictly speaking, a state without death, then by 'immortality' we mean an infinite existence. If something that exists is exempt from death, and death is non-existence, then it is necessarily so that this thing will never not exist in the future. Therefore, immortality is a model of existence that is not overly

[2] I am of course discussing only secular immortal existence for a human life and not any kind of highly speculative divine existence after death theorised by various religious doctrines.

[3] See appendix A for more discussion on what we mean or should mean by immortality.

relevant to discussion of life extension, because immortality means total exemption from death and it therefore not a realistic model of existence for humans. As a result, this thesis will be concerned with life extension rather than immortality, because, as will be obvious to many readers by now, any amount of life extension will always fall short of immortality. Even substantial life extension falls short by an *infinite* amount when placed next to a necessarily infinite existence, an existence rather difficult even to fathom.

Though some philosophers, for example, Leon Kass and Bernard Williams—whose arguments I examine in Chapters 3 and 4—use the term 'immortality' when discussing what might just as easily be called 'life extension', important and relevant points can be extracted from their arguments that can be applied equally to longer lives of the less than infinite kind without being committed to talking about immortality. Accordingly, this thesis will not be concerned with the rationality and permissibility of immortality as such, but rather the rationality and permissibility of life extension, for example, profound life extension of many centuries.

1.1.2 Extending the life of the organism or the life of the person?

Our goal is to extend human lives but what is it exactly we want to extend? Do we want to extend the life of the human organism, the biological living body, so that it survives for as long as possible? Or do we ultimately want to extend the life of the person, so that they survive for as long as possible? More importantly, which one is more valuable, human organism survival or personal survival?

I would suggest that those who desire life extension desire precisely that, personal survival, the extension of 'my' life. The survival of the organism as such is an irrelevant matter. I will, therefore, refer to survival of the person simply as 'survival'. Survival is all I should care about, if I survive as someone then I do not die. Survival, as I've been talking about it, might be said to presuppose identity. However, this might present a problem for some philosophers. We want to be able to define identity as a kind of 'non-branching' survival, but it is possible to 'survive' as two or more people, for example, if I were exactly replicated and the original me and my copy both survived (Parfit 1984). Nonetheless, analyses of problems of identity are beyond the scope of this

thesis and I will assume that survival as a result of life extension is survival of the same person.

So personal life extension simply means survival. Thus extending the life of the person is what matters if one cares about extending one's own life and the goal of extending the life of the organism is at best a means for the extension of the life of the person. What if, however, the extension of life eventually led to the survival of the organism but the non-survival of the person? This is to say, what if, as a result of certain limitations of human desire and character, further life extension resulted in the survival of the organism and the creation of serial lives or serial persons, only one of which is 'you' at the start of the series. This is one of Bernard Williams's worries about life extension, and the problem constitutes one half of a dilemma he suggests faces the individual who wants to live a profoundly extended life. He thinks that personal desire for life extension would be irrational if what one cared about was attractive personal survival, because either life extension would eventually result in the non-survival of the original person if one were to let one's desires and character change in order to keep an indefinitely long life novel and to avoid permanent boredom, or boredom would be the consequence of allowing one's desire and character to remain the same throughout. Either kind of life, Williams argues, would be unattractive and meaningless for the individual. Thus the desire for profound life extension is a self-defeating one.

Of course, it does not necessarily follow that life extension would be a bad thing even if we were to agree with Williams that to desire it is irrational, just that one would have no reason to extend their lifespan if their aim was personal immortality. (As discussed earlier, and irrational action may make one better off. Rationality and prudence may come apart). I will discuss Williams's fascinating arguments for the undesirability of profound personal survival, as well as counterarguments to Williams, in more detail later in the Introduction and at length in Chapter 3.

One more thing we can say about survival of the person and survival of the organism is that they have different values depending on our question. For example, if our question is whether personal life extension is a rational preference, then survival of the person is the primarily valuable thing. Whereas if our question is whether people living longer lives in the world is a good thing, then survival of the person is of neutral value at best, because the existence of human beings or bodies on the earth is what matters for our answer.

1.2 Life extending strategies

I begin Chapter 2 of the thesis by looking at the dominant theory of human ageing, the evolutionary theory of ageing, and what it implies for the potential success of life extending interventions of the anti-ageing kind. Some interpret the theory as strongly suggesting that human ageing cannot be manipulated, while others interpret the theory as suggesting that ageing and the maximum lifespan can be manipulated.[4] I do not argue in favour of one side or the other in the empirical debate but merely mark it as one of the ongoing hurdles to successful life extending interventions of the anti-ageing kind.

I proceed with Chapter 2 by outlining in detail some of the actual anti-ageing strategies that are being developed. The most promising strategies to retard or reverse ageing in humans include: (a) caloric restriction, which is the strategy of restricting food intake in order to produce effects conducive to the slowing of age-related diseases and decline; (b) the introduction of the enzyme telomerase into normal human cells, which is thought to increase cell lifespan; (c) gene therapy, specifically gene deletion, or the deletion of the gene 'SIR2', which is said to block anti-ageing interventions and thus is essentially an intervention to facilitate subsequent anti-ageing interventions; and (d) reversing age-related neural decline by inducing regrowth of lost synapses. Most of these strategies obviously go beyond direct intervention into the ageing process—if indeed it is a single process, which is also a highly controversial hypothesis in the empirical literature—but seek also to remedy the damage done by the ageing process on the human organism.

Alongside efforts to intervene into human ageing to extend the human lifespan there are other life extending strategies whose target has nothing to do with retarding or reversing human ageing and whose success will mean extended lives of a different kind. The first potential strategy is body part replacement. Work done on embryonic stem cells and human cloning may reveal the potential to replace ageing and worn body parts with healthy new parts grown artificially from human stem cells. What might the pursuit of this strategy look like? Consider, for example, the 'Borg' from *Star Trek*. The Borg are a species whose ultimate teleological end is the technological perfection of

[4] For the former interpretation, see Hayflick 1998; for the latter, see de Grey *et al* 2002a, and Kirkwood & Rose 1991.

themselves as well as other 'frail' biological races through the gradual replacement of inferior biological parts with superior technological parts. One by one the organism's faculties are 'enhanced' through their replacement with technology. A by-product of this quest for technological perfection is longer lifespans. The relentless adaptation of individual members of the species in order to make them resistant to damage or deterioration from ageing or external harm has made the Borg long-lived and on their way to becoming invulnerable.

An immediate reaction to such an aggressive, non-biological approach to extending human lives might be one of repugnance, that this course of action would make us more 'machine' than human and not simply in the physiological sense, but that our thinking and behaviour would change as we not only looked but felt more machine than human. But the strategy of body part replacement is at work today, to a less profound and somewhat more biological extent, when vital organs are transplanted into dying patients to replace injured or failing ones, extending their lives. Sometimes new parts are not always biological, but technological, for example prosthetic limbs. Body part replacement as a life extending strategy could potentially give humans limitless capacity to keep on replacing ageing parts of their bodies and brain, perhaps even wholesale replacement of the brain, and indefinitely hold off the biological warranty period for the human organism. In short, with continual body part replacement, humans could conceivably live forever as cyborgs.

Then there is cryonics. Taking inspiration again from speculative fiction, cryonics is a life extending measure that could allow a vastly extended, if discontinuous, survival. For cryonics to work, first the individual has to die, then their body, or simply their brain is immediately frozen with the aim of preserving it in as life-like a condition as possible. Then we wait, a decade, a century, until the appropriate level of biotechnological advances have been made so as to revive the brain and body and hopefully bring the person back to full life and health. The life extension permitted by cryonics amounts to a kind of time travel allowing the individual to escape causes of death now and survive to a time when death is truly avoidable. Consequently, the ultimate goal of cryonics is not simply to make it so the individual exists in the future but for individuals to be able to take advantage of future technology for curing all other human ailments. It is life extension as a kind of insurance policy, that upon resurrection in the future, one has waiting for them superior life extending interventions and restorative marvels (though, of course, this is not a certain bet).

The final non-ageing-related life extending strategy is by far the most speculative. What we do is back up the entire contents of your brain on to a digital storage device for the future resurrection of your consciousness or identity. This has been called 'mind uploading' or sometimes just 'whole brain emulation' (Astakov 2007). In the end, there is survival of the person but not of the organism. In fact, nowhere does the typical human form survive, unless the receptacle which houses the consciousness is shaped like a human being, perhaps for reasons of aesthetics or nostalgia.[5] For a fictional example, in the novel *Neuromancer* by William Gibson (1984), a character is able to interact with the 'saved consciousness' of a deceased computer hacker, who is accessible by 'jacking in' to an external storage device on which he has been uploaded. Similarly, in the science fiction TV series *Battlestar Galactica* (2004), a race of cybernetic intelligences called Cylons are able to upload the consciousness of a deceased Cylon into another identical body, a kind of technological metempsychosis, effectively making individual members immortal if repeated indefinitely.

Mind uploading could be another case of a discontinuous life extending strategy like cryonics, in the sense that you are at one moment living an embodied life, then 'you' are removed from the organism and 'rebooted' from a digital storage device containing your consciousness and your body is done away with. However, the interval between embodied and disembodied existence might be quite short, unlike the cryonics method. So having one's mind uploaded might really be no different from going to sleep one night and waking up the next morning as a consciousness on a hard drive.

Having described some of the major age related and non-ageing-related life extending strategies, there is a final dilemma about the desirability of either variety of life extension. The dilemma is this: either life extension is unlikely, because of the empirical uncertainty over whether we can in fact manipulate ageing and the lifespan this way, or life extension won't work, because the other biological and non-biological strategies—especially strategies like mind-uploading—would not ensure survival. Therefore the desirability of potential life extending strategies is problematic. We

[5] Maybe we can still say that the person or consciousness that survives is of a typical human form, but what does a disembodied consciousness look like? Perhaps nothing, because it gets its 'shape' or limitation in space in virtue of the body in which it is instantiated, whether this is a biological organism or a computer or something else. Think about the Internet, for example, which doesn't exactly exist or have shape in space except when it is made tangible via a computer. If there were no computers the Internet would cease to exist in space, but would be 'waiting' to be instantiated. Could the same, perhaps, be said about a human consciousness if there were no 'bodies' for it?

cannot be certain that direct intervention—retardation or reversal—of the ageing process will be possible. Meanwhile the other life extending strategies, though avoiding the ageing controversy, instead involve a threat to survival in virtue of the fact that they are strategies that involve dismantling or replacing the organism in order to preserve the person.

While I suggest in Chapter 2 that successful life extending strategies of the anti-ageing kind might not be available in the near future, it would be reckless to conclude that we will never be able to manipulate the maximum lifespan internally this way. Likewise, I will assume in the thesis that life extending strategies of the non-ageing-related kind will not present a problem for survival of identity. Thus the thesis will proceed from the assumption that we will have life extension one day in the future and that it will ensure 'your' survival. So the question is: Would life extension be a blessing or a curse?

1.3 Would life extension be a bad thing?

Chapter 3 of the thesis begins by examining two prominent sets of arguments that the desire for a profoundly longer life is irrational and that such a life would ultimately be a bad thing for the individual. Leon Kass (2004) offers a series of claims for the conclusion that a human life that went on indefinitely would be a meaningless life or detrimental to the goal of living a meaningful human existence. For instance, Kass claims that if we lived too long and were under little time pressure to engage in meaningful relationships or complete substantial projects, we would cease to treat these important human activities with the seriousness they deserved, making such a life less consequential. Kass is arguing that the value of meaningful human activities comes largely from our recognition that we only have a finite amount of time to engage in them and the fact that our lives are not infinite in time is essential for our human existence having meaning.

To say that individuals who potentially had endless time on their hands would be prone to live a less serious or meaningful life has some psychological basis. Take, for example, human beings' susceptibility to procrastination. Some people are more disposed than others to put off the completion of important activities or to seriously

engage with life even with the knowledge that their life will be limited. We usually feel that this lack of urgency or commitment makes for a poorer, less consequential life. Just as some people are at greater risk of procrastination given too much time, life extension may benefit some people while harming others. But Kass's worry seems to be more general. Given the nature of human desires and interests, if people had longer lives than they already do, procrastination and ambivalence would become universally prevalent and this would make human existence generally poorer.

We can reply to Kass by making a distinction between, on the one hand, having knowledge of how long one will live, and on the other, the actual length of the life itself. Imagine two hypothetical kinds of person. The first person has endless personal time; they are in fact immortal but believe their time is finite. Thus they are 'psychologically' mortal or have mortal attitudes to time. The second person believes they have endless personal time but it is in fact finite, i.e. mortal. Kass's concern about loss of seriousness and meaning seems to be a problem only for the second kind of person, the one who 'thinks' they are immortal and will have endless time to complete projects, participate in meaningful relationships, etc. Whereas the first person, who thinks they have but limited time to complete goals or to satisfy desires, even though they are in fact immortal, does not seem to be in danger of their life losing seriousness or meaning. What does this do to Kass's argument? It means that Kass's true target is not the actual fact of an endless personal existence but an impractical psychology or attitude to time exhibited by particular kinds of persons.

It's uncertain, however, how this distinction helps the individual who has immortality. We could not simply urge the true immortal to pretend or deceive themselves into thinking they will have a limited lifespan in order to avoid the threat of frivolity and meaninglessness. In any case, procrastination is a risk even within the confines of the relatively short lives we currently live. So procrastination seems like a problem even for a person living a finite life of, say, a few hundred years. They might become accustomed to putting things off for greater and greater lengths of time. However, even if this was necessarily a bad thing, as Kass wants to argue, it need not be something that destroys the value of the life entirely. Think about the kinds of strategies employed by most of us during our short lives to prevent procrastination. We have lifespans that are sufficiently long to make procrastination not only possible but also a frequent threat to accomplishment and engagement. However, we fight the yearning to delay and postpone by setting ourselves artificial goals whose completion requires less

effort and less time as we, bit by bit, work our way towards accomplishing the original goal. Couldn't we continue to employ this strategy given a much longer lifespan? It would amount to a similar thing as putting psychological limits on one's life, as suggested above, that is, thinking about one's life in a manageable, piecemeal fashion and looking forward only to the completion of one modest goal at a time.

Even if we agreed with Kass that the personal desire to live forever is irrational, it does not follow that life extension is a bad thing. Kass wants to say that not only are longer lives bad for individuals, but that it is bad that individuals live longer lives. Individuals would suffer, which would be bad, but the fact that there were longer lives might nevertheless be good, personally or impersonally. The only thing Kass's arguments potentially demonstrate is that it would be irrational for an individual to have the psychology of immortality, that is, to go through life thinking they are immortal and have unlimited personal time. So not only have Kass's arguments not shown that life extension is bad—that persons living longer lives is a bad thing—but they have also not shown that it is irrational for an individual to desire profound personal survival under the right kind of psychological constraints.

Bernard Williams (1973) presents an even deeper challenge to the personal desire for a profoundly longer life. He presents a sort of dilemma for the individual who is deciding whether to embark on an endless existence. Williams suggests that a profoundly longer life will play out in either of two ways: either the individual remains largely the same kind of person with the same desires and character or they allow these things to change over time. In the first of the two alternative lives, the individual will necessarily reach a state of permanent boredom, says Williams, because all meaningful desires and goals will eventually be exhausted, making continued life not worthwhile. In the second life, in which the individual allows their desires and goals to change over time, it will not be certain that the life of the person living in the remote future would be attractive anymore to the individual, that it would indeed be a different person living in the remote future to the one who started out on the life. Williams concludes that neither life courses would be desirable for the individual. Therefore, at the very least, desiring a profoundly longer life would be self-defeating or irrational.

Like Kass, Williams wants to generalise from his conclusions about the rationality of the desire for profound personal survival to claims of badness, but even if Williams is right that it would be irrational for an individual to desire profound personal survival, it does not follow that life extension is impersonally bad or that it is bad for

individuals to live longer lives. The two questions come apart as we saw at the start of the introduction. An answer to the question, 'Is it better for me to live a profoundly longer life?' does not constitute an answer to the question, 'Should I extend my own lifespan?' or, 'Should we as a society develop life extension?'

Let's look at the first horn of Williams's dilemma a little more closely. Would a much longer life, lived with more or less the same kinds of desires and goals, eventually result in a state of permanent boredom? If it can be demonstrated that this will happen, then it would be irrational for the individual to desire life extension. But again, the real question is not whether an extended life will result in a state of boredom, but whether this renders the entire life not even worth desiring. Williams seems to believe that an extended life must not include boredom for it to be worthwhile or attractive to the individual wholesale. Indeed he says at one point that "nothing less will do for eternity than something that makes boredom *unthinkable*" (1973: 95). There seems to be no further reason for Williams's insistence, besides the fact that a state of boredom would possibly involve some suffering for the individual. Maybe it is that Williams recognises that the nature of things is not in the budding immortals' favour. Echoing an Epicurean sentiment (which I discuss in detail in Chapter 4), Williams says "there is no desirable or significant property which life would have more of, or have more unqualifiedly, if we lasted for ever" (1973: 89). But even if 'all novelty is but oblivion', even if no model of an endless, satisfying state or activity, that would never prove boring or become meaningless over a profoundly extended life, can be found, this does not necessarily make the life that eventually ended in a state of permanent boredom meaningless or undesirable in its entirety. What about all the good parts leading up to the state of permanent boredom? Why must the individual who decides to start out on an extended life necessarily think the whole enterprise not worth pursuing because he knows that at the end of eternity he will inevitably experience boredom? Even if an extended life ended in boredom, maybe all the extra good things allowed by a longer life outweighs its badness, and makes the sacrifice worthwhile for the individual. Their desire for life extension might not, in the end, be irrational after all.

Another concern is this. Why must we assume, as Williams seems to, that boredom is a permanent state? This doesn't connect to our regular experience of boredom, which is a temporary state one falls into but then comes out of. It's simply not psychologically realistic that there could be such a state that, as soon as one reaches it, one necessarily becomes a permanent victim of it. Some have suspected Williams of an

unjustifiably jaundiced view of the limits of human desire and amusement. For instance, John Martin Fischer (1994) has suggested that while some pleasures are "self-exhausting"—these are pleasures derived from activities that, once obtained, we have no desire to repeat—we are also capable of experiencing "repeatable pleasures". Repeatable pleasures are derived from activities that no matter how many times they are repeated we experience no loss of desire for them or enjoyment in them, for example, good food and drink, sex, and other activities, like listening to music, viewing a sunset, etc. These repeatable sorts of pleasures are exactly the sorts of thing that could be used to fill out an extended life and keep boredom at bay indefinitely, or so Fischer argues.

That being said, Williams is ultimately talking about a state of permanent boredom, a state resulting from the necessary exhaustion of one's categorical desires, desires that are beyond the mere cravings of physiology or momentary pleasures but are the desires that make the life of a person meaningful or worthwhile, the goals that motivate a life forward into the future. When these kinds of desires or ambitions run out—perhaps because these are the kinds of desires Fischer calls self-exhausting—then continued life is meaningless for the person. Consequently, I think that Fischer has not yet given a sufficient criticism of the first horn of Williams's dilemma.

Let's now consider the second horn of Williams's dilemma. This can be interpreted as asking whether life extension would still be 'attractive' for an individual, and if not, whether this discounts the rationality of desiring it. If my desires and goals were to become so radically different from my current ones then a much longer life might not be an attractive prospect after all. Fischer interprets Williams's argument as one about the attractiveness of profound life extension to the individual. He accepts that one would 'survive' and suggests that the real puzzle presented by the second horn of Williams's dilemma is not about whether I can be certain that the desires and goals of the person living the extended life in the future will be recognisably 'my' desires and goals, but whether I am made better off by having different desires and goals as I live an extended life. Williams thinks that you are worse off because they will not be 'your' desires. But there is an obvious sense in which your desires three hundred years from now are your desires, just a different 'selection' of them out of the entire collection of desires that make up your extended life. So the question is whether the fact that different desires and goals instantiate themselves later on in your ongoing life makes you worse off. I would suggest that having different desires and goals over the course of

a radically longer lifespan is, at worst, neutral, and that it's hard to see that 'you' are made worse off by this.

Could you be said to be better off because life extension allowed you to have different desires and goals within the same life? Perhaps you're not made better off either. The mere addition of extra desires and goals to a life does not seem to make you better off than you would have been, presuming we're talking about the kinds of long term desires or goals that give a life meaning. It is perhaps of neutral value for the individual, by virtue of the fact that your desires and goals, at the time you are pursuing them, simply are the particular desires and goals you have at that time. To have different ones at a different time—assuming they're desires and goals that bring equal pleasure upon being satisfied to the previous ones you had—is simply to be trying something different, neither better or worse, like desiring strawberry ice-cream when you once desired vanilla ice-cream. But having the *opportunity to* pursue different desires and goals is, if anything, better than not having the opportunity. So life extension is at least not bad for you in terms of your interest in surviving individual life extension, and may, in the sense of expanded opportunity, be a good thing.

Williams's worries about the problem of meaninglessness for individuals living profoundly longer lives do not seem to be contingent, like Kass's worries, on what I've called the 'psychology of immortality', that they would only be concerns for the individual who went about life thinking they are immortal. Consider the individual who does not go through life thinking they have an infinite amount of time to do things, even though they may live an indefinitely long life. This individual could still get to a point in their life where they are 300 years old and are able to say that all of their desires and goals have been exhausted and they have reached a state of permanent boredom. So the worry about boredom is still a problem and is harder to dispel than the putative problem of meaningfulness, so too is the alternative problem of living serial lives in order to avoid boredom.

Although I remain undecided on the question of whether an individual's capacity to sustain indefinitely the kinds of desires that make a life worth continuing (what Williams calls categorical desires) over an indefinitely long life, I am satisfied that the alternative to living a life of fixed desires and goals would not present a problem for an individual, certainly not so much as to make the prospect of this second kind of extended life unattractive. Allowing one's desires and goals to change might be the only kind of life possible for individuals wanting to live profoundly longer lives,

and it's doubtful that Williams's claim that this kind of life would be in some way harmful should dissuade them looking forward to extending their lifespan.

1.4 Would life extension be a good thing?

I claim that it is rationally permissible to want to live indefinitely and that a longer life is not necessarily bad for us. But this is only half the story. There is a further question: Is more life a good thing for a person, ceteris paribus, so that we are rationally required to want to continue living? We need an answer to this question because we need to be able to present a reason why developing life extension—presumably by diverting funding and resources away from other human endeavours, technologies, or medicines—would be an important thing for humans to do.

The most obvious way to proceed is to show why death, the thing that prevents extension of life, is a harm to its victim. If death weren't in some way bad for the person who dies, then using resources to develop life extending technology to extend the human lifespan would probably be unjustified.

Not everyone has thought that death is a bad thing for the one who dies, or that having more life is any better than having less life. Epicurus famously presented an argument that death is never a harm to the one who dies, thus there is no reason to worry about it or avoid it. In his *Letter to Menoeceus* Epicurus urges the reader to understand this view of death:

> Death, therefore, the most awful of evils, is nothing to us, seeing that, when we are, death is not come, and, when death is come, we are not. It is nothing, then, either to the living or the dead, for with the living it is not and the dead exist no longer.

Presented formally Epicurus's argument seems to be this:

(1) If P is dead then P does not exist

(2) If P does not exist then P is not in a harmed state

(3) Therefore, if P is dead P is not in a harmed state

*(4) Therefore, death is not a harm

It might look to many readers like Epicurus is making a trivial point: Your death cannot harm you because when you are alive you are not dead and when you are dead you don't exist. We might readily accept this 'existence requirement'. But Epicurus seems to want to say that death is in no way a bad thing for the one who dies. This we are not so ready to accept. Most of us think that death is frequently a bad thing for the one who dies, for example, if the one who dies is very young.

We can concede to Epicurus that he is right that it is not bad for the person to be in a state of being dead. However, while this is true, it is still bad to *die*. Death is not 'the state of being dead', death—the thing that harms—is an 'event' in a person's life. The 'event of death' is that moment which takes one into the state of being dead, the event that takes a person from existing to not existing. When the person is dead, death has already happened. So death is an event that happens to a living person in that final moment before they are in a state of non-existence. Although I have formulated Epicurus's argument as referring to 'the state of being dead'—which the famous passage quoted above seems to imply—this is not *all* that death is. The state of being dead is not the death we are talking about when we say that death harms the one who dies. What we mean by death, then, the thing we want to say "harms" someone is an ante-mortem event, the very last ante-mortem event possible, because the subsequent state is being dead, with nothing in between. The living person is harmed in that final moment before they die by having any further opportunity for good experiences, preference satisfaction, etc., deprived of them. Thus it is bad to die, or to put it another way, the 'event' of death is bad for the one who dies. Therefore, in response to Epicurus's further inference, death is a harm for the one who dies: they are harmed in that final moment of existence in which everything is taken away from the person.

Epicurus says that death can only be a harm when you are alive or when you are dead, but it is not a harm in either. But Epicurus is wrong about the first horn of this claim, because death is a harm when you're alive. Can Epicurus answer this new challenge? If we reformulate Epicurus argument from above so that death is no longer 'the state of being dead' but the 'event of death' or the moment the person 'dies' we get this:

(1) If P dies then P does not exist.

(2) If P does not exist then P does not suffer harm.

(3) Therefore, if P dies then P does not suffer harm.

*(4) Therefore, death is not a harm.

But I do exist during the event of my death, which is a last moment of existence, the moment I die. I am harmed in this last moment of existence by being denied the opportunity for existing for more moments, for being denied a future. Therefore, premise (1) of the argument above is false. The subject being harmed is a living person; it is that living person who a moment ago had a life ahead of her, which has now been taken away.

The Epicurean conclusion about death, that it is in no way bad for the one who dies, seems unsatisfactory for most of us because we are able to reflect on the badness of being deprived of the goods of life, and also opportunity for more life, by the event of death. I support in Chapter 4 a version of the deprivation account of death's badness but where death is the 'event' of one's death, as described above. In this way, my version of the deprivation view manages to avoid certain counterintuitive implications and problems encountered by Thomas Nagel's (1979) original version of the deprivation account of death's badness. All of this I argue in more detail in Chapter 4.

Developing life extending technology to prolong the human lifespan is a worthwhile pursuit for the reason that death can be a bad thing for the one who dies, making postponing death, ceteris paribus, a good thing. My deprivation event view of the badness of death demonstrates this, while at the same time avoiding problems faced by Nagel's deprivation state view of the badness of being dead. Obviously not everybody will desire more life in cases where death is not a bad thing for the one who dies, but the deprivation event view of the badness of death allows for these cases. But by virtue of this account we are rationally obligated to prolong lifespans in cases of premature deaths.

1.5 Consequentialist duties to extend life

We are now in a better position to address the central normative question of the thesis. We want to know not whether survival is worthwhile for the individual but whether we, i.e. societies or governments, should develop technology to extend human lives. Would having life extending technology be beneficial or harmful for societies and future generations?

This is a very complicated question to answer. The development of life extension would have effects on both present and future generations of people, putting existing people in a position where we have to make a decision about what future populations would look like, for example, in terms of bringing fewer people into existence as a strategy to balance out existing people having longer lives. In Chapter 5 I analyse a rich argument presented by Peter Singer (1991) that it would be impermissible to bring about a world in which we have life extension.

Singer draws this conclusion about life extension based on a hypothetical scenario of the near future in which a life extending pill has been successfully developed that doubles the average lifespan from 75 to 150. He imagines some limitations to the pill's capabilities. The life extension pill only takes effect at the onset of middle age regardless of when the individual starts his or her routine, and though it grants the user temporary prevention of their ageing process (at least up until the very end of their life) the extra years permitted by the pill are lived at a level of health and functioning not quite as high as that enjoyed in their natural lifespan. Thus the user of the life extension pill makes a sacrifice: extra years for a slight dip in quality of life. There is also a global sacrifice to be made, says Singer, if we decide to go ahead and use the life extension pill. We must all decide not to have children as frequently as we do now, as this will be the only way to ensure overpopulation does not occur, an outcome that would be bad for everyone.

Singer suggests that our choice of whether we develop life extension ultimately depends on whether it would benefit or harm future generations. By applying utilitarian principles Singer concludes that developing life extension would result in an outcome

with less than optimal total and average wellbeing, therefore developing life extension would not be permissible. However, Singer can be seen as only providing an answer, a utilitarian answer, to the policy question, that is: what would utilitarian principles say about whether societies or governments developing life extension would be permissible? He does not provide an answer to whether life extension would actually benefit future individuals.

In any case, there is a discrepancy in what Singer claims utilitarianism says about life extension and what utilitarianism actually says about life extension. Singer does not apply all relevant interpretations of utilitarian principles to the question of about whether we should develop the life extension drug. One interpretation of the average principle, for example, says that it is permissible to bring about an outcome if and only if the average happiness per person in that outcome is no lower than in any alternative outcome. This principle I call the average happiness per person principle. The conclusion we get by applying this principle to the case of developing a hypothetical life extension drug is that it would indeed be permissible, because the resulting outcome would contain greater average happiness per person. What does this conclusion about life extension demonstrate? It shows that not only will our choice of developing a life extending pill result in more happiness per person for future societies, but developing the life extension pill will also benefit future people, by increasing their level of happiness because of their longer lifespans.

1.6 Life extension and justice

In Chapter 6 I continue to discuss this question: 'Should we develop life extension?' But now we want to know what should governments do about life extension. If life extension were developed, say in the form of a life extension pill like in Singer's scenario, then societies that had the life extension pill would face a special question of who should get access to the life extension pill. The intuitive answer is that everybody who wants to extend his or her lifespan should get access to the life extension pill. However, since we won't be able to provide everyone with the pill, we are obliged to answer a question of how we ought to distribute life extension. Could there be a just distribution of life extension and what would that distribution look like?

We applied different moral principles to the question of whether developing the life extension pill would be permissible, and we can do the same to help us come to an answer about the present question of justice. A theory of justice is basically a moral principle about what governments ought to do. We might suppose that life extension will start out life as a privately funded, developed and owned resource, available, at least in its early years, only to the wealthy minority. The first intuition we might have is that it doesn't matter whether everybody is allowed to have a life extension pill, it is simply a resource like other speciality or non-basic good and a free market system is fair: those who can afford expensive life extension it can buy and use it, those who can't miss out.

But a resource like a pill that restores your health and functioning later in life and extends your lifespan is something of great benefit to most people; moreover most people would want it. It is likely that, after a while, governments will step in to buy out and subsidise the life extension pill and presumably attempt to make it available to more people than it was previously available to. But there still remains the dilemma about the fairest way to distribute the life extension pill, because it will be a limited resource, and not everyone who wants it will be able to have it.

How should governments distribute the life extension pill? One front-runner for an appropriate distributive theory in this situation is the Rawlsian theory of justice. The Rawlsian principle about how to fairly distribute resources says that a basic resource should be distributed equally unless an unequal distribution improves the condition of the least advantaged. Who is the least advantaged with respect to life extension? I would suggest it is those in society who will be the worst off with respect to their potential life expectancy. The worst off with respect to life expectancy includes only those groups whom the life extension pill will benefit. This excludes those who will die of diseases or who current cannot get their basic healthcare needs met. The life extension pill will be of little use to these groups. Thus my interpretation of Rawls's distributive principle becomes this: the life extension pill should be distributed equally unless giving priority to those least advantaged with respect to life expectancy makes them better off. This will be my answer to the question of how governments should distribute the hypothetical life extension pill.

1.7 Nonconsequentialist duties to extend life

Chapter 5 is a utilitarian consequentialist response to life extension. But what would a nonconsequentialist say about whether we should develop life extension? I examine these nonconsequentialist responses in Chapter 7. I begin with a general Kantian approach. On the Kantian view, what is most valuable is our rational nature, the survival of autonomous agents. Upholding our rational nature, our autonomy as rational agents, unlike the consequentialist approach, does not mean 'maximising' autonomy, i.e. that we ought to bring new people into existence who will eventually become full moral agents thus promoting the greatest possible amount of autonomy in the world. Rather it means respecting and prolonging the existence of already existing rational agents if this is their preference. Because maximising total or average welfare is no longer the obligating force from a non-utilitarian approach, it does not matter morally that world A (in which we don't use life extension) contains a greater total amount of welfare than world B (in which we use life extension), because there are more important values at stake, such as respecting the autonomous decisions of the people in world B, where we choose to use the life extension pill, to prolong their lifespan and uphold their right to continued life. The value of longevity on the Kantian approach lies in the value of the continued life of the already existing moral agent.

Therefore I would suggest that the Kantian would be in favour of developing a life extension drug for the capacity it would give people to fulfil their rational preferences. And since we have already addressed the question of whether it would be irrational, other things being equal, for the individual to extend their lifespan, this is not a problem for the Kantian.

It might be argued that a person has a conditional positive right to life, and thus a conditional right to life extending technology, on a Kantian view. One reason why a person's positive right, not just not to be killed, but to go on living is a conditional rather than a categorical one is that, since his value as a person is based on the value of autonomy of being a human agent, a value which is collectively possessed by other human agents, his value is his value in virtue of being a member of the group 'human

moral agents' and so is a value larger than any one person. Consequently the right to life of an individual person may sometimes be overridden by the right to life of other persons based on their own autonomous nature in certain circumstances. What all this amounts to, it seems, is that the rational nature of a person upon which the right to life is based is merely a feature that gives us a moral reason to support a claim to a right to life in certain circumstances, but makes it a right that is conditional on the similar right of other persons.

As to a person's conditional right to technology to extend their life on a Kantian view, this is again drawn from the Kantian principle for respect for autonomy. One principal way to respect a person's autonomy is to respect their desire to go on living, which can be done by entitling them to a right to life in the positive sense. We see this when we apply the right to life in the context of the older person's desire to go on living. So long as the older person is in possession of a right to life, in the non-consequentialist sense of an entitlement granted to autonomous agents, such a right to life in the very least should allow the possessor to choose to continue living as an autonomous agent, if this is their decision as it bears on their own life. If this is true, then in order to respect a person's autonomous decision to prolong her life we must provide her with the means to do so, that is, the medical and technological resources.

I would include a qualification of the Kantian conclusion about the permissibility of extending one's lifespan. I would suggest that we are talking about two different kinds of obligation in the case of a right to the technology to prolong the lifespan, other things being equal, and a moral obligation to prolong the lifespan. On a Kantian view, the former is a conditional right that those entitled may exercise or discard as they choose, and may clash with similar positive right to life of other persons under particular moral approaches. The latter implies an obligation to always prolong the lifespan regardless of individual desires, in a similar vein to those who argue that we are always obligated to keep people alive, a claim one might reject even if one believed that it were permissible to extend the lifespan.

1.8 Should we extend the lives of the young or the old?

In my last chapter I raise a further question of justice: When in competition for resources, are we obligated to extend the lives of the young over the old? Moreover, who deserves to extend their lives more, the young person or the old person? The answer might seem obvious: when in competition, the young deserve extra years of life because the old have already enjoyed more years of life than the young. This is essentially what is known as the 'fair innings' claim: that we should give priority to younger people over older people when in competition for life prolonging resources. The fair innings claim is a difficult one to reject, because of its natural appeal to our intuitive sense of justice. But we can still ask the normative question: is it permissible for governments to extend the lives of the young over the old?

It makes sense that those whose job it is to distribute viable life extending interventions—whether it be a pill like in Singer's story or something else—will not often face the dilemma of having to choose between giving the life extension pill to an older person at the cost of depriving a younger person of continued life. But this situation might arise in extreme cases. Then we must ask ourselves another moral question, a comparative one: Should we give priority to younger people over older people, or vice versa, when in direct competition for life extension? This is essentially the same question that we ask today when societies have to decide how to ration life prolonging healthcare resources, e.g. policy questions about whether more funding ought to be given to diseases of the old and less to other age groups.

There are arguments in favour of giving priority to younger people over older people, a prominent one is known as the 'fair innings' claim. Usually offered as a justification for rationing healthcare in favour of the young patient over the old one, the fair innings claim appeals to the strong intuition that it is only fair that everyone be given equal chance to reach a full number of years or a 'fair innings', and those who have already reached this threshold are less entitled to continued life then those who have not yet reached a fair innings. And there are also arguments that object to the bias in favour of the young over the old when continued life is at stake. These are so-called anti-ageist arguments. Both kinds of argument make claims about the moral significance of age. My task is to see whether these claims about the moral significance of age still hold or whether they are made less compelling given the possibility of

radically extending the human lifespan, which could involve the retardation or reversal of ageing.

I conclude in this chapter that the fair innings intuition is unjustly discriminatory towards older persons, since a plausible theory of distribution of life extension can sometimes give priority to old persons. There is no justifiable in principle reason that older persons should get less priority for life extending technology than younger persons.

CHAPTER 2

CAN WE EXTEND HUMAN LIVES?

> We have the means right now to live long enough to live forever. Existing knowledge can be aggressively applied to dramatically slow down aging processes so we can still be in vital health when the more radical life extending therapies from biotechnology and nanotechnology become available.
>
> Ray Kurzweil, *The Singularity Is Near* 2005: 371.

The latest empirical evidence is less than conclusive about whether we will one day be able to extend human lives. There is little consensus among biologists and ageing researchers as to whether, among other things, human ageing can be successfully manipulated. The major controversies in the empirical literature about potential life extending strategies include: whether the dominant theory of why humans age, the evolutionary theory, leaves open the possibility of manipulating ageing; whether human ageing is controlled by a single process or by multiple mechanisms; and whether there is a biological limit to the lifespan of the human organism.

In light of these disputes I allege that some theorists have been overly optimistic about the likelihood of radically extending the human lifespan in the near future, at least via interventions into ageing. A sensible weighing of the evidence suggests that the successful application of life extending interventions on humans is not likely to occur within the next twenty years as some have predicted.[6]

This chapter has evolved out of the need to deal with the existing empirical arguments surrounding human ageing and longevity as things that we can manipulate. Furthermore, there is the need for a reliable empirical base from which to inform the ethical issues about whether we should develop life extension that I will analyse in the following chapters. But I will not be presenting an argument for or against the permissibility of developing life extension in this chapter. What I will present is a discussion of the major debates among scientists about whether the human lifespan is

[6] Biogerontologist Aubrey de Grey and colleagues have previously stated that within a decade (from 2002) we could possibly reverse ageing in laboratory mice. de Grey *et al* 2002a. de Grey has also stated in an article that once ageing has been reversed in mice, we could reverse ageing in humans by the following decade. de Grey 2004.

something that can or cannot be extended. This will involve discussion of theories of human ageing and longevity, prominent anti-ageing strategies, along with other more speculative kinds of life extending technologies and their possibility of future success.

Despite my somewhat negative conclusions in this chapter regarding the timescales for the development of successful life extending interventions, in the chapters to follow, I argue that significantly prolonging the human lifespan would indeed be permissible, and seek to rebut various objections to developing life extending technology.

2.1 The debate over human ageing and maximum lifespan

Most of the literature on life extension refers to strategies for retarding or reversing adult human ageing. Therefore, in order to gauge the success of proposed anti-ageing and other life extending interventions, we must first consider some of the prominent theories about why humans age. The theories expressed in the following discussion fall into two categories: those that leave open the technical possibility of manipulating human ageing, and those that do not.

2.1.1 The evolutionary theory of ageing and its interpretations

Evolutionary theory implies that post-reproductive anti-ageing strategies confer no survival value for a species and consequently these strategies are not selected for. This has suggested to some biologists that human ageing and the maximum human lifespan cannot be easily genetically manipulated. This is indeed the opinion of Leonard Hayflick, who gives a sobering evolutionary account for why humans experience ageing (1998: 643):

> Humans…had a life expectation at birth of 30 years or less for more than 99.9% of the time that we have inhabited this planet. Prehistoric human remains have never revealed individuals older than about 50 years of age. There appears to be no selective advantage favoring the survival of old animals or old humans…Indeed, biological aging may have arisen as a phenomenon coincident with the appearance of the human species. The only animals that experience significant age changes are humans or the animals that we choose to protect. Because humans, unlike feral animals, have learned how to escape the causes of death long after reproductive success, we have revealed a process that, teleologically, was

> never intended for us to experience. One might conclude, therefore, that aging is an artifact of civilization.

Before the arrival of a protective civilisation a premature death for humans was highly likely, due to hostile environmental factors such as predation, starvation, and infectious diseases. The phenomenon of adult human ageing had not yet revealed itself. Faced with the prospect of a short lifespan due to the dangers and instability of a wild existence, organisms discovered an effective means for getting around the problem of mortality: have reproduction take precedence over the individual organism living longer.[7] Genes were passed on as early as possible once reproductive maturity had been reached, making longevity a redundant capacity from the 'viewpoint' of nature. According to the evolutionary account, then, humans experience ageing because mutations enhancing survival post-reproduction have no selective advantage. We never needed them. It is difficult to see how evolution could have selected for the characteristic of longevity when it is likely few organisms would have lived long enough to take part in the selection process (Hayflick 1998). Humans, says Colin Farrelly, might, therefore, lament that "our biological design is driven by the shortsighted aim of continuity of the species", that "once our reproductive capacities have run their course there is no evolutionary benefit to keeping a particular person in their prime physical and mental capacities" (2010: 390).

The implications of the evolutionary theory of why humans experience ageing have given rise to two opposing interpretations. Either, the evolutionary theory of ageing suggests that adult human ageing cannot be manipulated and that any limits to the lifespan are fixed. Or the theory clearly leaves room for the manipulation of ageing and extension of the limits to the human lifespan, both of which are currently suboptimal. Hayflick is the leading proponent of the first camp. The evolutionary story suggests to him that humans have a biological warranty period that limits the lives of individuals and the life expectancy of populations. We humans are defaulted to wear out and intervening in ageing requires more than the removal of our ageing mechanisms. Recently demographers have used the latest evidence of patterns of morbidity, i.e. disease, and mortality for humans to suggest the possibility that human bodies are limited by just such a biological warranty period, a term the demographers apparently use as a more accurate substitute for 'longevity' or 'lifespan' (Carnes *et al* 2003).

[7] This is not to imply that reproduction is an evolved trait, just that the existing possibility of reproduction was biologically favoured.

Hayflick himself calls the idea that an organism's longevity is determined 'longevity determination' and likens it to something comparable to an engineer's notion of the 'mean time to failure' of any mechanical device (1998: 641):

> …after reproductive success an animal has the potential to survive for a period of time determined by the level of excess physiological capacity reached at sexual maturation…that is what I call longevity determination…The level of physiological capacity reached at the time of sexual maturation determines the potential for continued longevity.

Essentially, what Hayflick is saying is that the time after reproductive success for an organism is a period of "coasting", where developmental processes have ended and the organism's capacity to maintain its vital systems declines. Just how long this 'coasting period' is, and the 'level of excess physiological capacity' reached, determines the continued longevity of the organism. But it is still a period of longevity that, from the point of view of nature, has redundant value for the continuation of the species. As a result, the general attitude of those scientists who assert that human ageing and lifespan cannot be manipulated is that nature knew what it was doing in making us age and die after a limited lifespan, and it is futile to think to do what evolution cannot or did not need to do.

On the other hand, scientists such as Aubrey de Grey (*et al* 2002a; 2003), Tom Kirkwood (1991; 2002), and Michael Rose (1991), believe that the evolutionary theory does in fact leave open the possibility of further prolongation of the human lifespan. They have pointed out that what natural selection actually does is select for the optimum rate of ageing for each species given its evolutionary niche. Thus *non-ageing* beyond this rate is always thought to be "suboptimal", or below or less than the optimum level that it could be. The evolutionarily-determined optimum lifespan of a particular population does not necessarily have to be the maximum lifespan. Thus they argue evolutionary theory strongly suggests that, in fact, organisms are programmed for survival not death, and that human longevity may 'evolve' in the future, along with our slowly changing evolutionary niche, to give our species greater longevity (Kirkwood 2002).

So according to the Kirkwood and Rose view, an organism's longevity can be manipulated. Evolutionary theory itself seems to be compatible with this view, because different species have acquired different distributions of lifespan, with humans apparently being one of the longest lived species, in fact, the longest lived mammal.

Therefore it seems premature to conclude that the human lifespan is fixed with no chance for manipulation. This response might even suggest that since the environmental milieu of many human populations is now less hazardous, natural selection might slowly produce or evolve a 'new' species with a much longer lifespan even without the need for the genetic manipulation of ageing. Of course, such an evolution of human longevity would require many future generations to see any change, whereas those who want a longer lifespan want it now, in their current lifetime.

Hayflick and Kirkwood and Rose all seem to agree on the basic evolutionary account of *why* humans experience ageing—namely, natural selection favours mutations that enhance fitness early in life, but the force of natural selection in eliminating a mutation such as ageing is weaker after a person's reproductive years allowing adult ageing to persist—they simply disagree about what it *implies* as far as being able to manipulate ageing and the lifespan.

On the one hand there is Hayflick who interprets the evolutionary story as saying that curing ageing requires more than the removal of our ageing mechanisms, even that curing ageing is probably not possible, but is rather a necessary function of the human organism. All of which leads to the conclusion that there is a fixed limit to the maximum lifespan of human beings. Kirkwood, Rose, and de Grey, on the other hand, interpret the evolutionary account of ageing to mean that ageing is not necessarily absolute, all the story shows is that the force of natural selection in preferring survival to longevity is stronger directly after reproductive success but weaker in later age, and that acquiring longevity. So the possibility of non-senescence or non-ageing is always an option, and in humans is something that is always operating at a level less than what it optimally could be.

I do not have an argument in favour of one interpretation of the evolutionary theory of ageing over the other, but rather would like to highlight that this disagreement is a contributing factor in my conclusion to this chapter, that researchers have been too optimistic about developing successful anti-ageing related life extending technology any time soon.

2.1.2 Does human ageing involve multiple mechanisms or a single process?

The evolutionary account of ageing suggests that human longevity and senescence are under the control of heritable characteristics, and that the origins and processes of this complex adaptation involve random variables, making the prospect of successfully manipulating human ageing all the more unlikely. Thus, theorists who favour the evolutionary explanation of human ageing and longevity not only advance a theory about why humans live as long as we do but that multiple mechanisms are involved in controlling much of human ageing and that there is no single 'ageing process' that may be easily manipulated.

Those who believe human ageing involves multiple mechanisms have often been broadly referred to as "complificationists", whereas those who believe that there are but a few or perhaps one single major mechanism controlling ageing have been referred to as "simplificationists". [8] For the simplificationists, the possibility of developing successful interventions to manipulate human ageing seems much more likely. The simplificationist also rejects the hypothesis that evolutionary theory suggests that humans have a fixed maximum lifespan.

The simplificationists get their inspiration from observations of gene mutations in the likes of rodents and roundworms, which have resulted in resistance to age-related disease and a prolonged lifespan (Flurkey *et al* 2001; Martin 1999). These results have suggested to experimental ageing researchers that there could indeed be a single mechanism that may control maximum lifespan: you manipulate one factor and lifespan is extended. Similarly, there are the results gathered from the technique of caloric restriction, or (CR), which is the strategy of limiting an organism's dietary intake in order to lower cholesterol and blood pressure, and as a result, slow ageing.

Evidence from caloric restriction studies on rodents and a variety of other short lived species that successfully generated life extending effects also suggests a single mechanism controlling ageing, at least in short-lived non-human species. If the simple manipulation of one environmental factor such as the restriction of food can have a marked effect on ageing then perhaps ageing is only a matter of a few or even a single mechanism after all. CR's success, however, has only so far been replicable in a variety of non-human species, not in humans, and therefore it remains to be seen whether CR's

[8] This distinction is found in Martin 1999.

success in lower-order animals supports a case for a single mechanism controlling human ageing. What these experimental observations suggest is that control of lifespan and cell senescence could indeed come from simple genetic manipulations.

Despite the breakthroughs with CR and observations of gene-mutations in non-human animals leading to longer lifespans, theorists such as George Martin (1999: 18) still conclude that ageing is under "highly polygenic controls", referring to the fact that it is determined by a number of genes, favouring the complificationist view of human ageing.

2.1.3 The concepts of 'lifespan' and 'life expectancy'

There are two important concepts to grasp in any debate over human longevity and these are the concepts of maximum 'lifespan' and average 'life expectancy'. Maximum lifespan refers to the outer limits of human longevity, the maximum number of years that a human can (currently) live for.[9] The concept of maximum lifespan is central to the contemporary debate. Some theorists like Hayflick (1998; 2000), as we have seen, maintain that this maximal value is fixed at about 125 years for evolutionary reasons.

Life expectancy, on the other hand, refers to the average age of death, given the current environmental conditions. Thus, lifespan can be seen as something biologically determined, whereas life expectancy presupposes this biological notion of a limit to the lifespan and is something like a statistical conception of the mean number of years an individual can expect to live a healthy, functioning life given their environmental milieu. Over the past hundred years average life expectancy has been steadily increasing in modern Western societies.

Two demographic researchers have made the case that throughout history the best extrapolations, from the then current data regarding life expectancy limits, have been repeatedly broken by an average of five years after publication of those life expectancy calculations (Oeppen & Vaupel 2002). They argue that this evidence, along with the trend of increasing life expectancy, supports the case that average life expectancy will continue to increase, at a rate of about a quarter of a year per year, indefinitely, and that there is, in principle, no limit to human life expectancy. They point

[9] Trotter 2004, esp. pp. 198-200, about how elusive the idea of 'lifespan' is within the medical community.

out that if this life expectancy trend continues, average life expectancy for humans will reach 100 years within the next six decades, and this means a baby born in the 2060s can expect to live to *at least* 100 years of age. There is also agreement among other demographic researchers that, regardless of the deliberate pursuit of life extending technologies, the current trend of increasing life expectancy is likely to continue, with the most realistic estimate being that humans can expect to live around 130 years, on average, by the year 2100 (Richel 2003).

However, some demographers and some evolutionists disagree, because this contradicts the 'maximal lifespan of humans' theory of some biologists. If the average life expectancy of humans were to continue to increase for an indefinite period, then it should eventually, on paper at least, surpass the putative 125-year mark. However, this is not taking into account that the factors that contribute to the trend of increasing life expectancy—such as better health care, medicine, nutrition, education, etc—are factors *extrinsic* to the biology of ageing and therefore may have no ultimate effect on the maximum human lifespan which is related to the failure of the organism due to ageing.

Making the distinction here between intrinsic and extrinsic influences on longevity further underscores the different goals of anti-ageing researchers I discussed in Chapter 1. When the goal is to intervene in ageing to directly increase maximum lifespan, then the focus is solely on the manipulation of intrinsic factors, i.e. the ageing process itself. When the goal is not to directly intervene in ageing, but rather to compress morbidity, the aim might be the manipulation of extrinsic factors, which may indirectly have an affect on life expectancy, but may have no affect on pushing the boundaries of maximum lifespan.

We have seen that the concept of the human lifespan refers to the maximum number of years a human currently lives, whereas life expectancy refers to the average age of death given current mortality conditions. The variations in life expectancy between populations are due to environmental factors, factors extrinsic to the biological ageing of the organism, like life style, nutrition, medical care, etc. However, it has become increasingly evident that a growing number of the world's populations today are living to ages where the lifespan appears limited only by the failures of the biological organism associated with ageing. These failures of biological functioning are intrinsic to the biology of the organism.

However, Gerald Gruman (1966) has suggested that the ill-defined concept of 'old age' has allowed for the idea of the 'absolute lifespan' to remain almost

unchallenged and those whom it would benefit to have established the notion of a fixed lifespan. Moreover, others have found, based on experience from the autopsy table, at least up to the 1960s, that we never see a death from pure 'old age' alone (most older persons die from age related diseases), casting further doubt on the existence of a natural limit to lifespan. The implication of our traditional conception of old age, in Gruman's words, is that "the life-span phenomenon might be an artifact of the statisticians and not a never-yielding barrier" (1966: 7).

The environmental factors that contribute to increased life expectancy may have no impact in pushing the boundaries of the lifespan of humans. In spite of continuing increases in life expectancy, there could still be a biological limit to lifespan because of failures within the organism caused by ageing. In any case, it is increasingly clear to anti-ageing researchers that the only way to increase the lifespan of humans would be via the manipulation of the mechanisms of ageing, if indeed these mechanisms can be isolated.

It is clear from these empirical debates, there is no consensus yet as to whether the lifespan is something that can be extended beyond its current maximum, a threshold that may or may not have been pre-ordained by the evolutionary processes we humans have experienced. These empirical issues are, however, complex. The characteristics of a species, including maximum lifespan, are going to be a function of both biology and environment. This is still compatible with there being a biological limit to the lifespan, however, a biological limit to lifespan is not fixed for a species forever, and there may always be room for manipulation, as suggested in section 2.1.1.

Therefore, the interesting question is whether the environmental factors that have been found to contribute to increased life expectancy over the past century might contribute to our maximum lifespan in small ways in the future. Some suggestions regarding this could be that: (1) environmental factors dispose an individual organism or population to increased longevity over a range of generations, or (2) they dispose an organism or population to the possibility of a mutation it otherwise would not have experienced. There may be evidence to suggest that either of these outcomes may eventuate. We may certainly find that human beings, in our evolutionary niche, may gradually evolve to experience greater longevity in the remote future.

Now that we have set up a theoretical framework for understanding some of the empirical controversies involved in the goal of prolonging the human lifespan, the

remainder of the chapter will focus on setting out what are the latest attempts at lengthening the human lifespan, and setting out the best estimates for success.

2.2 Promising anti-ageing life extending strategies

Since the 1950s, it has been one of the primary goals of biogerontological research to develop interventions to slow and, if possible, reverse ageing. Contemporary ageing researchers refer to any strategy whose aim it is to achieve an increase in the lifespan of an organism by combating biological ageing as 'strategies for engineered negligible senescence' or SENS, where the term 'negligible senescence' refers to the *absence* of ageing in an organism (de Grey *et al* 2002a: 453).

These anti-ageing strategies, though still in the experimental stages, are mostly untried on humans. They are primarily carried out on laboratory mice and other shorter-lived species because their life and health extending effects are achievable far more quickly than would be possible in humans. While some researchers maintain that a negligible senescencing (non-ageing) human is a daydream (Hayflick 2000; Masoro 1996), there is still no indisputable evidence that attempts to *engineer* a non-ageing mammal are always going to be unsuccessful. While the intervention of only a few of the above strategies might not achieve the end of restoring youthful vigour and reversing the mechanisms of ageing in humans, the "coordinated implementation" of them *all*, say de Grey and colleagues for example, could indeed do so (2002a: 460).

The following are some of the most promising interventions to slow or reverse ageing according to the empirical literature.

2.2.1 Caloric restriction and slowing ageing in humans

It has been known since the 1930s that restricting the food intake of laboratory mice makes them live longer (McCay *et al* 1935). This intervention has since come to be called caloric restriction, or CR, and involves restricting caloric (energy) intake in laboratory subjects, while still maintaining sufficient quantities of vitamins and minerals and other required nutrients, in order to lower indices like cholesterol and

blood pressure, both biological indicators of ageing. It has been reported that CR has been observed to extend what could be called the 'health span' of rodents, with CR rodents remaining healthy and vigorous at ages where the control rodents have long since died. Such CR rodents exhibit much greater exercise capacity, and autopsies show much lower levels of arthritic (inflammation and swelling of the joints), neoplastic (tumours), and other degenerative change (Miller 2002).

However, in spite of almost 70 years of research into the life extending effects of CR on a variety of species it is still unknown whether CR increases lifespan in humans (though, to be accurate, the human-specific question has only been directly addressed for the last 20 years). CR studies on non-human primates, like the rhesus and squirrel monkey, suggest a cross-species effect of CR on ageing, and therefore its similar life extending effects in humans are not unlikely (Heilbronn & Ravussin 2003; Lane *et al* 2002; Roth *et al* 2001).

CR researchers as yet cannot quantify the length of lifespan extension humans could potentially enjoy due to CR intervention; nor does the literature tell us whether CR, in order to produce its maximum life extending effects, would need to be administered earlier or later in human life. The one thing that the evidence has consistently indicated, however, according to Roth *et al* is that CR is the *only* anti-ageing intervention to have been shown "conclusively and reproducibly" to slow ageing and delay age-related disease in mammals (2001: 306). This explains the continued attention given to CR as a real hope for intervention in human ageing.

An explanation for the lack of empirical evidence of the effects of CR in humans is to some extent owing to the great difficulties in conducting long term studies of CR on free-living human beings.[10] Even in the event that CR is found to have the similar effects of slowing ageing and extending lifespan in humans as it does in other species, the strict dietary control regime involved in CR would be unrealistic for most humans to maintain (Lane *et al* 2002). It has therefore become important to researchers that an alternative strategy to the strict dietary restrictions be developed if CR is to have a realistic potential as a life extending intervention for humans. One such alternative is called CR 'mimetics', interventions that 'mimic' the protective effects of CR without resorting to the strict regulation of food intake (Roth *et al* 2001).

[10] An exception is a recent study by Stein *et al* 2012, which found significant differences in the age-associated changes in autonomic function between a group of 32-80 year olds on a CR diet compared to the control group on a regular Western diet, with the heart rate variability (HRV) of CR individual similar to those individuals 20 years younger eating regular Western diets.

One such CR mimetic is "metabolic inhibition" at the cellular level. This intervention, say Roth and colleagues, may have the same effects of CR such as "reduced insulin and body temperature and neuroprotection" but without the need for the extreme dietary restrictions (2001: 337). Another possibility of a CR mimetic is a pill that mimics the physiological effects of eating less and dropping the large amount of calorie intake each day but without compelling people to go hungry.[11]

2.2.2 Telomere activation and cell life extension

Biologists believe that the finite capacity of the cells of an organism to divide is directly related to cell ageing of the organism post-reproduction. In humans, normal human cells apparently undergo a limited number of cell divisions before they eventually reach a state where they can no longer replicate contributing to age changes in the organism. There have been recent encouraging reports that researchers have been able to genetically intervene in cell senescence and kick-start the replicating process once more and circumvent this cellular longevity determination through the introduction of an enzyme called telomerase into normal human cells (Bodnar *et al* 1998).

The tips of chromosomes (chromosomes are the structured form of DNA in cells) were discovered to be discrete structures called telomeres, and the thinning or lessening of telomere length can explain the limited capacity for cells to replicate (Greider & Blackburn 1985). This shortening of telomere length in human cells apparently 'signals' the cells to enter senescence, or cellular ageing. In any event, researchers took this phenomenon to be ageing at the cellular level, and cell senescence is therefore considered to be an 'end-replicating problem' (Hayflick 1998).

In the 1960's, following experiments on cell cultures, researchers observed, along with the class of normal cells, a particular kind of 'immortal' cancer cell that could replicate indefinitely. Following this experimental observation on cell cultures the process of immortalization of normal cells has naturally drawn much attention in the field of cytogerontology. The question researchers have asked themselves is how normal human cells could attain similar unlimited capacity to replicate as that found in some cancer cells.

[11] The question of whether such a pill could be developed was asked recently by Lane *et al* 2002. The serious search for an anti-ageing pill, *Scientific American*, 287: 24-29.

In the mid to late 1980s researchers Greider and Blackburn (1985) identified an enzyme called 'telomerase' that was found to add "repeats" to the ends of chromosomes and maintain telomere length. The introduction of the enzyme telomerase into normal human cells is thought to increase cell lifespan. This intervention looks to be at least as potentially successful as CR intervention in slowing or reversing ageing, but as Carol Greider, one of the identifiers of the enzyme telomerase, has reported, it is a discovery that "raises more questions than it answers" and will require much more work on such empirical questions related to "how telomere length might signal entry into senescence", and whether or not humans "may have evolved to have additional mechanisms to protect against cell immortalization and cancer" (1998: R180).

2.2.3 Rejuvenating lost synapses and prolonging cognitive functioning

As the neurons of most brains are known to be significantly worn-out with age the possible regrowth of synapses could extend the longevity of the human brain and cognitive functioning. The successful treatment of cognitive impairment in aged laboratory rats by administering human nerve growth factor to reverse age-related neural decline has already been reported in the literature (Chen *et al* 1995).

There is the additional possibility of the replacement of brain cells. Certain neurological conditions involve the loss of cells from the brain and nervous system through disease (e.g. Parkinson's and Alzheimer's disease) or injury (e.g. traumatic brain injury) that are not naturally replaced and to which researchers have suggested cell replacement therapies to replace the depleted 'cell populations'. This is done by the work of neural stem cells which are cultured in an artificial environment outside of the living organism, and which, following injection into the brain, have been stimulated to recognise and replace lost neurons (Armstrong & Svendsen 2000).

The *reversal* of age-related neurodegeneration is thought by researchers to be more promising than the *replacement* of lost brain cells engineered through cell replacement therapies (de Grey *et al* 2002b). However, some theorists may worry about the fact that the second strategy of replacing cells in a person's brain might threaten their survival. The philosopher's concept of survival is different to the general notion of survival, which is usually taken as the continued functioning of the biological organism, the survival of the body or human being. Survival for philosophers doesn't necessarily

entail survival of the organism, but is more accurately survival of the person, the 'me' or the thing that is important in personal survival. Though cell replacement therapy does not involve the 'wholesale' replacement of the age-impaired brain (as some more speculative strategies, discussed later, might propose) it does involve a gradual replacement of impaired or lost brain cells. However, it would be quite unmerited to suggest that the new healthy 'version' of the person's brain is so radically different that the person themselves has changed.

The strategy, it seems to me, is in the spectrum of possible strategies that involve a gradual process of replacement, which is a process the human body itself performs on a regular basis when it continually destroys and rejuvenates cells within the body. Derek Parfit discusses an "imaginary" case of radical brain surgery that, similarly, involves the gradual replacement of small parts of the person's brain over the course of about a hundred operations.[12]

The issue he raises is whether, after the operation, this procedure is as good as normal survival for the person concerned. If by replacement of small parts of the brain Parfit means something like individual cells, then I agree that this procedure, on the face of it, would be indistinguishable from the 'natural' replacement of brain cells which has been observed to take place regularly in the human body with apparently no danger to the survival of identity.

Issues of survival and identity are often raised in conjunction with the potential effects of new technologies on persons and life extending technologies are no exception. I will, in the next section, touch on issues of personal survival in relation to life extending technologies of the more speculative kind. There are genuine problems of personal identity raised when considering the possible life extending techniques of body-part replacement, and mind-uploading (discussed in the next section). However, to the extent to which such questions of identity are explicitly raised in the discussion in the present and in any future chapters, I will regard the solving of these identity problems as beyond the scope of my current project and instead assume that any life extending interventions that are eventually developed preserve identity.

[12] See: Parfit 1984. 'Appendix D, 'Nagel's Brain'', pp. 474-77.

2.3 Other life extending strategies

2.3.1 Body-part replacement

There are, of course, many speculative strategies for extending the human lifespan. The outer limits of science being able to provide us with ways of staving off death and keeping the human brain and body alive and functioning are highlighted by such science-fiction-esque interventions. One suggestion has been that further research into embryonic stem cells and human cloning may reveal the possibility of replacing ageing body parts with 'new' healthy parts grown artificially from human embryonic stem cells (Keller & Snodgrass 1999).

This intervention could conceivably give humans a limitless capacity to keep on replacing the ageing parts of their body and brain, perhaps even the whole brain, and indefinitely extend their lifespan. This strategy, like the strategy of cryonics discussed below, is different from the above anti-ageing interventions as it does not strive to understand human *ageing* for the purpose of slowing or reversing it, but rather, the body-part replacement strategy involves prolonging the 'warranty period' of the human brain and body, to keep it functioning by implants or attachments of healthy replacement 'parts'. With continual body-part replacement, the human being could conceivably live an eternal life as a 'cyborg'.[13]

Of course, there may or may not be misgivings about the replacement of body parts, but in the event of our brain being replaced wholesale, aside from it being a potentially impossible task,[14] it would not be clear, assuming some version of physicalism, whether the same person would indeed continue if the brain were 'traded-in' in this way. Granted, this replacement-of-brain procedure would most likely involve a kind of precise cell-for-cell replacement, essentially an identical replica or duplicate of

[13] A famous thought experiment regarding the persistence of identity is that of the 'Ship of Theseus' which asks the question whether an object (or person) whose parts (or cells perhaps) were constantly replaced with new parts would fundamentally remain the same object (or person). Refer also back to the discussion of the potential one-by-one replacement of cells and neurons in the brain to restore cognitive functioning.

[14] Olshansky *et al* 2002, call this "more the subject of science fiction than likely science fact".

my original brain. Nevertheless, it is the possible wholesale replacement of the brain that presents a problem for some theorists who believe we are essentially our brain.[15]

For example, imagine a future society where wholesale brain replacement was the only successful way to extend one's maximum lifespan. In this society it becomes a choice between having my original brain replaced with a duplicate, and therefore having a prolonged lifespan, or not having the brain replacement procedure and living a shorter lifespan. Two important questions arise for the person who takes the former option with respect to her continued existence: (1) When my brain is copied would *I* be copied, that is, would the result of the operation facilitate the continuation of my identity? (2) If receiving an identical replica of my original brain involves no 'subjective' difference, i.e. all my memories, desires, traits of character, etc., are continuous, does it even matter?

It has already been suggested above that a more successful strategy would be a gradual replacement of brain cells, at a rate that poses no danger of compromising the identity of the subject. But even in this scenario of successive cell replacement or rejuvenation within the brain, one may argue that a finite number of replacements of brain cells could lead to a qualitatively different brain and thus a different person.

It seems in many ways that the limits of our longevity, and we may also say of our identity, become apparent when we consider the replacement of our brain in the same way as we think of the replacement of other old body parts.

2.3.2 Cryonics

There is another, perhaps even more speculative, strategy to prolong the human lifespan: cryonics. The strategy behind cryonics as a 'life extending' measure is that first the individual must die. According to Robert Ettinger, cryonics and general transhumanist pioneer, he or she is then immediately frozen in such a way as to preserve their body "in as nearly life-like a condition as possible", and while in suspended animation, awaits the appropriate level of biotechnological advances "equal to the task of reviving and curing" their brain and body (especially the brain) back to full life and health (1965: 1-2). Obviously, while in suspended animation for an undefined amount

[15] Thomas Nagel is apparently one such theorist, see: Parfit, D. *Reasons and Persons*, 'Appendix D, 'Nagel's Brain', pp. 468-77.

of years, decades, or even centuries, the individual is not continuing to live a life during this time. It is a discontinuous existence, with the possibility of more life in the individual's future. This is how life is 'extended' through cryonics.

Cryonics is about avoiding *permanent* death, rather than extending the lifespan conventionally, by betting on the efficacy of future biotechnological and medical developments. This makes cryonics different from other potential interventions, which aim to extend the human lifespan in a way that allows the individual a continuous life. In these other strategies, such as the anti-ageing strategies discussed above, the individual would continue to live the same way they had been only hopefully for a longer and healthier period of time.

Most of us are familiar with the basic technique of cryonics from modern science fiction stories. Indeed, the speculative nature of cryonics as a life prolonging intervention is due to its origins in speculative fiction. For example, in the novel *Neuromancer* (1984), by William Gibson, a wealthy, aristocratic family who live in a compound-like residence on a luxury space-habitat, have for decades been extensively using cryonics technology to prolong the lifespan of each generation, of mother and father, sons and daughters. The patriarch of the family is apparently 'thawed out' every 30 years or so, and is currently 200 years old. Similarly, in the world of the novel *Ubik* (1969), by Philip K. Dick, one's deceased spouse or relative can now be accommodated, provided you have the money, in a moratorium facility using cryonics technology which sustains them in a life-after-death existence, and, with the use of a special drug called 'Ubik', allows them some level of communication with the outside world, e.g. with the 'bereaved' who come to visit them in their state of suspension.

It must be said that the goal of cryonics is, again, not in-principle impossible despite its great practical difficulties, there is nothing completely incoherent about the system of thought behind it as a life extending measure, but it is, nonetheless, at least to most contemporary people, a very unlikely scenario and a technological feat perhaps beyond our ability now or in the future. Critics, such as Michael Shermer, disparage the unrealistic possibility of being reanimated after frozen in suspended animation and point out the flaws in such thinking. In particular the immense difficulty of 'reanimating' the frozen brain of a human being is singled out (2001: 29):

> …thaw out a can of frozen strawberries. During freezing, the water within each cell expands, crystallizes, and ruptures the cell membranes. When defrosted, all the intracellular goo oozes out, turning your strawberries into runny mush. This is your brain on cryonics.

However, cryonicists, not to back down without a fight, turn to the potential of emergent nanotechnological wonders to solve this problem. They believe that nanotechnology will facilitate reanimation by allowing microscopic 'machines' to be inserted by injection into the defrosting patient, at which time these machines can repair the billions, perhaps trillions, of molecules and cells of the brain, restoring them. Whether deserved or not, sceptics refer to cryonics as "borderlands science", and this is because, as Shermer again articulates: "it dwells in that fuzzy region of claims that have yet to pass any tests but have some basis, however remote, in reality…It is not impossible for cryonics to succeed; it is just exceptionally unlikely".

This theoretical possibility of cryonics technology successfully reanimating a frozen and dead human being allows the experimental cryonics researcher the privilege of keeping intact the dream of cryonics as still within the reach of our future technological powers.

2.3.3 Mind uploading

Could we one day back-up the entire contents of a person's brain like we can do the contents of our computer hard-drive and then store such information on removable devices for the future resurrection of that person's consciousness? An analysis of the hypothetical resurrection of the human consciousness, or 'mind uploading' (sometimes called 'whole brain emulation'), has already received some attention (Astakov 2007).

The possibility of copying, scanning, and uploading one's consciousness into a more durable surrogate 'body', be it a storage device or even a 'virtual body', has increasingly become the plot device for many futurist science fiction stories in recent years and will continue to be a prominent idea in the future for experimental researchers as well. For example, again in the novel *Neuromancer*, the protagonist is able to interact throughout the story with the 'saved consciousness' of a deceased but legendary hacker by way of a 'ROM-module' in order to use the hacker's expertise on a job. The entire contents of the hacker's identity was apparently saved onto the external storage device (much like today's CD or DVD-ROM's we assume, but with enormous storage capacity!) that can be loaded-up and 'jacked into' in order to access the stored 'data' and interact with the personality of the deceased person.

What is so astounding about this hypothetical longevity strategy is that once the contents of a person's brain is copied and uploaded onto some super computational device their immortality is more or less assured, the information won't degrade or deteriorate, at least not for many centuries, when it could conceivably be copied once again to a different external device.[16]

Mind uploading is perhaps a misguided hypothesis. As an *a priori* possibility it may be called into doubt when it comes to what it assumes our consciousness is and its relation to the physical world. To begin with, mind uploading carries the assumption of a materialist view of ourselves and of the universe. The most common form of materialist response is to suppose that the mind is simply the brain, that the entirety of our inner lives, our sensations, emotions, perceptions, and memories, can simply be reduced to complex physical processes in the brain. This kind of reductive materialism is still a somewhat controversial hypothesis, which some who still hold out against it, suggest is an all-too-easy answer to the problem of the nature of the 'mind'.[17]

This being said, the idea is that the mind's empirical manifestation is a contingent matter, the brain is one such manifestation, but the mind can be captured by a functional definition, it is the function that matters and that is why other physical realisers, like storage devices, might succeed in prolonging human consciousness.

In any event, the issue at stake regarding the hypothesis that we could one day upload the contents of a person's brain onto a supercomputer appears to be the issue of whether this technique might preserve identity. What would make it so that *I* survived a mind-upload intact? Could we ever, like Parfit's teletransporter, copy the exact states of all the cells in a person's brain and transfer this information into another body or another kind of storage device, effectively creating an exact copy of the person's brain and thus facilitating the continuation of their personal identity along with it?[18]

Thought experiments like this show once again perhaps only the theoretical possibility of the copying and transferring the contents of an individual's brain. But for now it is a hypothesis more at home within the realms of the science-fiction novel, like Gibson's *Neuromancer*, than as a practical intervention for prolonging the human lifespan.

[16] Significantly, by the end of Gibson's novel, the hacker wants to be "erased" from the ROM-module and his existence permanently ended, to which he eventually gets his wish.
[17] For an argument against a reductive explanation of consciousness, see Chalmers 1996.
[18] Parfit 1984, especially pp. 199-201.

2.4 Conclusions

I believe it is safe to say that humans are not going to radically extend the maximum lifespan any time soon, though incremental increases to life expectancy will likely happen due to those extrinsic factors I discussed earlier. The "within ten years for mice and twenty years for humans" estimate for the reversal of ageing given by de Grey and colleagues is overly optimistic given a review of the empirical literature.

The main contributing factors for delay in successful life extending interventions are the lack of consensus among theorists as to the correct interpretation of human longevity theories, the shortage of successful trials of anti-ageing interventions on humans, and the lingering biological questions still outstanding for researchers as to exactly why interventions like telomere activation and CR work as they do. Although de Grey and others may eventually be proven right regarding their timescale for the reversal of human ageing, the empirical evidence just doesn't support such claims at this time.

Though we are ignorant of when successful life extending interventions will become available it is quite obvious that life extension will become available in the future and the consequences of longevity research will have great significance for the lives of individuals and human societies collectively.

Notwithstanding the empirical uncertainties discussed above, the goal itself of radically prolonging human lifespans asks us to address some general questions about whether doing this is self-evidently a good thing. Is there benefit to be gained from the traditional limits to the human lifespan? Is it a good thing to avoid death? How might the development of longevity technology influence the lives of future generations?[19] These are some of the ethical and related philosophical issues I will attempt to address in the following chapters.

[19] Walter Glannon, for instance, argues that the manipulation of ageing—e.g. genetic interventions like telomere activation, caloric restriction, etc—could negatively effect future generations of people by altering their lives due to what he believes will be a "higher incidence of mutations not selected against in earlier life" Glannon 2002: 341. This is just one possible genetic consequence of anti-ageing interventions for future generations.

CHAPTER 3

WOULD LIFE EXTENSION BE A BAD THING?

> …death is not necessarily an evil, and not just in the sense in which almost everybody would agree to that, where death provides an end to great suffering, but in the more intimate sense that it can be a good thing not to live too long.
>
> Bernard Williams, 'The Makropulos Case' 1973: 83

While I have suggested in the previous chapter that we will not have technology to radically extend our lifespans any time soon, I will assume now, and for the remainder of the thesis, that life extension is possible. Should we be hesitant about living considerably longer lives? Or would survival be such that we should gladly prolong it forever?

Many have indeed thought that there would be a problem with living too long. This anxiety has been encouraged since our earliest stories.[20] The theme invariably seems to be that endless personal survival would in the end be undesirable. The desire for endless survival in these stories is used almost as a teaching tool, an instruction and a caution against something one might end up regretting. Our usual intuition about all this is that, while we cannot know for sure that living an indefinitely long human life would eventually lead to a state of boredom, or some other kind of suffering for which death would be a release, we do infer from our own lives—from the existing structures of and limits to human pleasure and desire—that living too long might eventually become disagreeable. However, this inference is in tension with another strong intuition we have, namely, that death can be a bad thing because it makes our lives finite in duration.[21] I am assuming for the moment that this is an intuition that we all share, that death can be a harm. In the next chapter I will analyse this claim in detail.

---

[20] Early examples of the mixed blessing that is endless personal life include: the epic of Gilgamesh, an ancient Sumerian myth, in which Gilgamesh is envious that only the gods are able to live forever, but who eventually learns wisdom in the inevitability of death; the Greek myth of Tithonus, who, granted eternal life by Zeus, soon comes to regret his endless existence; Odysseus rejecting Calypso's offer of immortality in Homer's *Odyssey*; and the more contemporary tale of the immortal undead Count Dracula in his state of isolation and ennui.

[21] This is tension is evident, for example, in Bernard Williams 1973.

So we don't wish to die but we don't want to live forever either. We have two strong intuitions about our lives that are in tension: that death can be a bad thing because it prevents the opportunity for more life, and that too much life can be a bad thing because it will eventually run up against suffering. In this chapter I will examine the second intuition, that too much life can be a bad thing for the individual, making the desire for a profoundly longer life an irrational one. I will examine this intuition through an analysis of the most prominent contemporary arguments that have been given for its justification, those of Leon Kass and Bernard Williams.

I said that we make inferences about the potential badness of living too long based on contingencies of our experience of living the life of a human being, for example, how we experience pleasurable or interesting activities, and whether such activities could be indefinitely repeatable or whether our enjoyment of them must eventually become exhausted. In the case of an immortal life we point to contingencies like whether such a life would be lived in health and prosperity, or whether there would be other immortals to share the life with so it would not be lived in isolation. In this chapter I will assume that an endless life would be an endless existence as an embodied human being, that it would be lived in normal health and economic stability, and that it would be part of an outcome where at least some others are enjoying immortality alongside oneself. What we want to know, then, is whether profound personal survival, with reasonable opportunities for good health, wealth, and meaningful relationships, would be something rational to want and worthwhile to have.

The two writers whose arguments I have chosen to focus on, Kass and Williams, will give different kinds of reasons why endless personal existence would be a bad thing for the individual. In particular, Kass presents contingent reasons why personal immortality might be seen as a bad thing, whereas Williams argues that personal immortality would necessarily be a bad thing.

The reader will notice I have been using the terms 'immortality' and 'an endless life' to denote the target of these authors' arguments. When it is asked in the philosophical literature (and in everyday conversation) whether life extension would be a good thing the question is often put as whether 'living forever' or 'immortality' would be desirable. And this is how Kass and Williams discuss life extension. But I will examine their arguments as arguments against the desirability of life extension and in my discussion I will use the phrase 'profound personal survival' to more accurately capture the thing that Kass and Williams's want to show as undesirable. Even if it were

true that an endless life would at some point involve some loss of meaning for a person, this is not necessarily true for a lifespan that had been extended a few hundred years, which is still a long time for a human individual to live for. Therefore it is less clear, not to mention more interesting to consider, whether a life extended a few hundred years would necessarily be meaningless for an individual and not worth pursuing. For those interested in the concept of immortality or a never-ending human life I discuss basic assumptions about the concept of immortality and make some philosophical distinctions about immortality and human existence in Appendix A.

Personal life extension remains a strong desire of many and we generally think that a longer life, other things equal, is better than a shorter one. So it would be useful to know whether there is a point at which we can necessarily say that an individual has extended their life too much, has lived too long a human life. How much life is too much would, of course, be contingent on differences in individual stamina and interest. Nonetheless, some writers, most famously Bernard Williams (1973), have attempted to argue that too much personal survival would be intolerable for everyone and that death in this context would be a good thing.

What both Kass and Williams argue in the end is that it would be irrational to desire profound life extension. But even if it were irrational to extend one's life this does not necessarily mean that it would not be good for the individual to go ahead and do. Think about the not unfamiliar cases of individuals who, after sustaining severe injuries following an accident, are in a condition of unbearable pain and wish treatment to cease so that they may die, only to be forcibly treated by doctors and eventually recovering enough to live a worthwhile life.[22] It might have been rational at the time for these individuals to want to die, but we can say that cutting their own lives short would have been bad for them, because they went on to have lives worth living. So the question of whether it would be rational to want profound life extension comes apart from the question of whether having profound life extension would be a good thing and prudent for the individual to pursue.

[22] See for example, the case of Dax Cowart, who was severely burnt after an accident, which also killed his father, resulting in the loss of both his hands, eyes, and ears, and the loss of skin for over 65% of his body. While in hospital Cowart requested medical treatment to be stopped and that he be allowed to expire. The doctor's refused, forcibly treating Cowart for months, who eventually healed enough to leave hospital, went on to get a law degree and now has his own private practice.

3.1 Is our mortality good for us?

Some writers have tried to make the case that the things that give a human life meaning would be lost if we were to find a way to avoid a natural death and live much longer than we already do.[23] What these writers seem to have in mind is that there is some benefit to the individual, be it meaning or importance, to be drawn from a natural human lifespan and its termination in death. Leon Kass (2004), for example, has argued that individuals who want to avoid death or so radically prolong their lives are misguided because they do not realise the valuable things they will lose by pursuing a longer life. They do not realise, says Kass (2004: 311) that, "the finitude of human life is a blessing for every human individual, whether he knows it or not." Our mortality is what makes a human life meaningful and the fact that we all live a certain amount of time and then die has some force or weight. Thus an immortal life would be a meaningless one for a human person, argues Kass, and as a result not worth pursuing.

Kass has four separate arguments for why an immortal life would be a meaningless (and therefore undesirable) one for an individual. I will look at these arguments in turn and suggest why none of them ought to persuade the individual that profound personal survival would be a bad thing.

3.1.1 Loss of interest and the threat of boredom

The first argument is that, if we had too long a life we would eventually lose interest and engagement with it to the point of it being undesirable to continue living. New kinds of pleasures would not present themselves regardless of how long we lived and the same pleasures would not be sufficient to sustain us. Therefore a much longer life would be redundant from the perspective of the individual living it. The repetition of similar experiences seems to be the primary culprit. Kass (2004: 312) asks hypothetically whether the pro tennis player would really enjoy playing 25 percent more tennis games, or whether the Don Juans of the world "would feel better for having

[23] See: Callahan 1995; Kass 2004. Interestingly, certain religious thought argues the opposite, namely, that immortality, or belief in some kind of survival after physical death, is a requirement for life to have meaning. See: Metz 2003; Perrett 1986.

seduced 1250 women rather than 1000". The short answer to Kass's rhetorical questions is: Yes, they would. Both of these individuals, if given the extra years of healthy life, may very likely enjoy doing more of the activities they love to participate in in life. This extra pleasure would have value for them just as it did during their shorter lifespan. Kass is perhaps overlooking the fact that we don't normally live our life by mindlessly repeating a single activity until it has been done to death, if it is something that we truly enjoy and appreciate then we rationally spread out the activity so that it is repeatable and the pleasure we get out of it renewed every time.

Thus it seems that we could conceivably make the opposite inference to what Kass makes. I think many people when thinking about their own lives would confidently say that most of the activities that they engage in would only benefit from more time in their lives to pursue and enjoy them, only increasing their fruitfulness and their meaning. I am thinking here about such pursuits as close friendships, extended learning and intellectual pursuits, both of which would be greatly benefited by more time to develop and take pleasure in. Kass is perhaps again ignoring the extent to which our pursuits could be open-ended, that there can be various ways in which to do the activity that adds variety and new meaning and value conferred to these activities at different times during a life. There is also the obvious point to make that memory of the activity plays an important role in our repeated pleasure of it. Unless we have an infallible memory and that everything we've ever experienced is at our immediate recall, then for most of us our memory of an experience or activity fades over time so that when we participate in it again in the future the activity or experience seems fresh again and there is a renewal of pleasure. Why can't we consciously try to this in regard to our favourite activities indefinitely?

Kass derives inspiration from Lucretius, who Kass interprets as saying that, "We move and ever spend our lives amid the same things, and not by any length of life is any new pleasure hammered out" (2004: 313). The claim of this passage from Lucretius is that no matter how long our lives may be extended no new kinds of pleasures or activities could come about to sustain us throughout a longer lifespan and that we would simply spend all our lives essentially doing the same things over and over again. The passage articulates what Kass wants to say, namely, that an increase in life years does not come with an equal increase in the kinds of pleasures or goods of life, and this is one reason why an extended or infinite lifespan would be undesirable. But Kass is not allowing for the extent to which our activities could be more open ended and that we

may have desires that are not merely contingent on satisfaction of pleasures but are desires to achieve goals, desires that propel us into the future and give us reason to want a longer life in order to complete them. These kinds of desires, which we can assume for the moment that individuals have (Williams argues that we do, calling them "categorical desires"), need not necessarily be exhausted and lead to unbearable boredom for the individual leading an profoundly longer life.[24]

3.1.2 Loss of seriousness and commitment

Kass's second argument is that if we had extended lifespans we would have too much time on our hands, which would lead us to put off completing certain valuable human activities or committing to certain valuable relationships, and we would not treat them with the gravity or seriousness they deserved. As a result, these once important human pursuits are rendered less meaningful. The value of these human activities, says Kass, comes mostly from our recognition that we only have a finite amount of time to take pleasure in them. Thus if we had lives that were radically longer than what they currently are, these activities that make up a human life would lose much of their significance, to the detriment of the significance of the life that included them.

Kass asks (2004: 313): "Is not the limit of our time the ground of our taking life seriously and living it passionately?" He recognises that some human activities might not require the motivation that can only be given by our knowing we will die and only have limited time, and gives such examples as the human desire for understanding, and the best sorts of friendship. But he thinks these are rare exceptions, and claims that "most activities, and for most of us, I think it is crucial that we recognise and feel the force of not having world enough and time" (2004: 313).

As evidence, Kass appeals to fictional examples of the frivolity of the lives of the immortal gods of Ancient Greek literature. The authors of this literature thought that such beings would envy the finite life of mortal men because they are limited by death and have aspirations and life goals and genuinely feel the weight of their actions and those of other mortals likewise constrained by knowledge of their limited time alive. Man's mortality, says Kass, makes his life matter. But is Kass's argument really an

[24] There will be more discussion on open-ended activities and renewable categorical desires in the next section where I discuss Williams and his interpretation of the problem of desiring profound life extension.

argument against actually having an immortal life, or is it rather an argument that we should not *believe* ourselves to be immortal. This belief, that one will potentially live on indefinitely, by itself seems to be enough to give rise to the problem Kass is concerned with. Would the individual who did not believe she was immortal, whether she was in actual fact or not, be vulnerable to Kass's problem of seriousness and motivation? I suggest not.

So there are two questions worth highlighting about Kass's second argument in favour of mortality. The first question is, even if one were leading an immortal life, would one lose motivation to complete projects, work at relationships, and reach personal goals, if one knew that these things could be worked on at any time in the future? The second question is whether an immortal life would become less serious because of its actual limitless structure or because of the immortal's knowledge of its potentially limitless structure? I will examine each of these questions in turn.

Kass's claim that individuals who had too much time on their hands would be prone to a less serious or meaningful life has some psychological basis. Take, for example, human beings' proneness to procrastination. Some people are very prone to put off the completion of important activities or serious engagement with life, and we usually feel this lack of urgency or commitment makes for a poorer, less consequential life. Not only does the procrastinator seem to treat their own desires and goals less seriously, but those of others as well. So Kass's worry might be that, if people had even longer lives than they already do, procrastination and ambivalence would become more prevalent. And just as some people are at greater risk of procrastination if given too much time, so when it comes to significantly extending individual lives, life extension may benefit some people while harming others. Those well equipped to deal with the threat of procrastination may only benefit from having more time to complete valuable projects, while those not so well equipped to stay meaningfully engaged with life, might find that the liberty afforded by a much longer lifespan exacerbates their aimlessness.

Procrastination might indeed be a problem even for a person with a life of a few hundred years, that they might habituate unhurriedness, become accustomed to putting things off for greater and greater lengths of time. But is it a worse problem than that which we, ourselves face? Even if this was necessarily a bad thing, as Kass argues, might we not take example from strategies employed by most of us during even our modest lives to prevent procrastination? We have lifespans that are sufficiently long to make procrastination not only possible at the moment but a frequent threat to

accomplishing and engaging in things. But we fight this yearning to delay and postpone by setting ourselves artificial goals whose completion requires less effort and less time as we, bit by bit, work our way towards accomplishing the original goal. Couldn't we continue to employ this strategy given a much longer life? It would amount to a similar thing as putting psychological limits on one's life, that I suggested above, i.e. thinking about one's life in a manageable, piecemeal fashion, looking forward only to the completion of one modest goal at a time. Commitment and engagement will not get 'worse' in and of themselves given a longer lifespan, nothing about the nature of these things will fundamentally change, there will simply be more opportunities to fall prey to them throughout the course of a longer lifespan.

Kass is also clearly making a questionable empirical assumption about individual habits and behaviour, i.e. he is saying that commitment and engagement will get worse, a claim for which he needs more empirical support. This sort of thing likely differs between individuals who continue to take advantage of whatever extra years they have as vital and meaningful and not just as a collection of less vital or less important years to be wasted away. I am sure we can think of individuals we know who, fully aware that this is the only short life they get, still exhibit little motivation to pursue meaningful activities or complete valuable projects, etc., whereas others, with the same knowledge of their mortality continue to be inspired and motivated to investigate by every facet of life and to start new extended projects even in old age. People become motivated to various degrees and at various different times throughout their lives whether they consciously have in mind their limited lifespan or imminent death continuously feeding them the thought that they do not have all the time in the world in which to participate in activities or complete projects.

This leads us to our second question, that of whether an immortal life would become less serious because it was potentially endless or because the immortal individual believed it would be potentially endless. Could an individual ever know for certain that they would have endless time with which to participate in life? The contingently immortal individual (recall my definition of contingent immortality for Chapter 1) could not be confident that she will never die. I would suggest that the potentially immortal individual must always be contingently immortal in her own mind. She could not know that no external event will ever conspire to end her existence or that the universe itself will always exist to accommodate her. Thus she could never know

that her life will in fact go on forever. Nor could she have this belief either, unless she was deceiving herself or was optimistic to the point of irrationality.

I think, then, that we can make a distinction between personal experience of a potentially endless life and the fact of endless existence. As a result of this we can make the further distinction between two types of individual, one who in fact has immortality but who thinks they are mortal, and another, who thinks they are immortal but in fact will have a mortal life. Kass's worry about a loss to seriousness and motivation seems to be a problem only for the second kind of individual who thinks they will have endless personal time, but they need not be mortal in actual fact, but could in fact have genuine immortality. The problem would still arise for them solely because they think they will have unlimited time.

The distinction between an individual who is immortal but thinks they are mortal and one who is mortal but thinks they're immortal can be taken even further by considering the concept of 'personal time' as distinct from 'external time'. Personal time is time experienced relative to oneself, and is not really time like external time is, but has that role for the person. David Lewis first made the distinction between the notions of 'personal time' and 'external time' within the context of time travel in his paper 'The paradoxes of time travel'.[25] In that paper, he explains how the two senses of time come apart after illustrating the case of a hypothetical time traveller, which led to the question of "How can it be that the same two events, [the time traveller's] departure and his arrival, are separated by two unequal amounts of time?" (1976: 146):

> I asked how it could be that the same two events were separated by two unequal amounts of time...I reply by distinguishing time itself, external time as I shall also call it, from the personal time of a particular time traveler: roughly, that which is measured by his wristwatch. His journey takes an hour of his personal time, let us say; his wristwatch reads an hour later at arrival than at departure. But the arrival is more than an hour after the departure in external time, if he travels toward the future; or the arrival is before the departure in external time...if he travels toward the past....Instead of an operational definition, we need a functional definition of personal time: it is that which occupies a certain role in the pattern of events that comprise the time traveler's life. If you take the stages of a common person, they manifest certain regularities with respect to external time....The assignment of coordinates that yields this match is the time traveler's personal time....It isn't really time, but it plays the role in his life that time plays in the life of a common person...We can say without contradiction, as the time traveller prepares to set out, "Soon he will be in the past." We mean that a stage of him is slightly later in his personal time, but much earlier in external time, than the stage of him that is present as we say the sentence.

[25] I am grateful to one of my examiners for reminding me of the original source of this distinction and for urging greater clarification in my explanation of it.

So personal time isn't the same as external time (or simply time) but we can talk about it like it is a kind of time. I cannot do any better than Lewis in illustrating the distinction between personal time and external time, however, I might add that I would define personal time as time measured idiosyncratically by the individual—thus 'my' personal time can be different from 'your' personal time, but there is no universal, objectively measurable personal time, because that would simply be external time.

Both Lewis's and my distinction between personal time and external time might still seem unsatisfactory and difficult to pin down, but this is a testament to how difficult it is to explain the distinction. It might be that it is simply one of those fundamental distinctions, which we invoke when we need, but that giving an in-depth analysis of is very challenging. As I go on to explain in the thesis, it might be worthwhile thinking about the distinction between personal time and external time in terms of the distinction between a person's subjective experience of their life, and the fact of the life itself. Indeed this was the purpose of bringing up the personal time/external time distinction in chapter 3, namely, to show how there could be a person who thinks they have infinite personal time but has only finite external time. Or put another way, this person is not, in fact, immortal but experiences her life as if she were—or you might say, her experience could go on forever regardless of the external distance she has to travel.

Personal time is then perhaps equivalent to the person's experience of their life. This can be distinguished from the fact of the life itself. Thus there could be a person who thinks they have infinite personal time but has only finite external time, in the sense that this person is not in fact immortal but experiences her life as if she were, or you might say her experience could go on forever regardless of the external distance she has to travel. On the other hand, there could be a person who has finite personal time but who (for arguments' sake) in fact has an infinite external existence.[26]

Why is this distinction relevant? Think about how life extension is valuable. It is valuable because it extends your experience of life, i.e. your personal experience lasts longer, but it is not necessarily because life extension extends the external life itself. For example, when one is in the grip of a vivid dream, as far as their personal time goes, they may have experienced days, maybe even weeks, of meaningful events, when in

[26] Roy Sorenson (2005) was the first to employ this distinction in the context of immortality and life extension.

actual fact their external life has been one of lying inert and asleep in their bed for just a few hours.

One interesting thing to notice about the person with has infinite personal time is that they would not require a long life per se. This revelation might seem to work in Kass's favour, as this sort of person would not need actual immortality. But Kass's point is about an individuals' experience of their immortality rather than the fact of the external immortal life itself, so his worries would apply to the pseudo-immortal too, he who does not actually have external immortality, but because his personal time is potentially infinite.

Our finitude and our mortality can be separated. Likewise, I assume, our infinitude and our immortality. It seems that arguments like Kass's against the desirability of personal immortality do not account for how these things break apart. That one's personal time (or one's self experience) is finite doesn't entail that one is, in fact, mortal, and that one's personal time is infinite doesn't entail that one is immortal.

3.1.3 Dulling of aesthetic appreciation

Kass's third argument is that a life that avoided mortality, that knew too little of death, would eventually make the person unable to appreciate beautiful things; nor could such an individual genuinely love another. The meaning we give to the aesthetic aspects of human life is inextricably linked to an appreciation of death and decay and of ageing and decline, says Kass. He wants to impress on us the good sense that appreciation of beautiful and worthwhile things, "the beauty of spring warblers" or "the spreading sunset", and also the ability to create such beautiful things, as a piece of music or a poem about human frailty, depends on the knowledge of their impermanence and that only those who experience such impermanence themselves can create and find value in such things (Kass 2004: 314).

It seems quite clear to me that the impermanence of a beautiful thing is compatible with the observer's immortality, that the beautiful thing itself does not cease to be impermanent and thus lose its beauty even if the individual appreciating the thing is immortal. Therefore the conclusion does not follow from the premises in Kass's argument. Surely I could still appreciate beautiful things and consider them beautiful, or

love someone deeply even if I am living an immortal life. It is unclear why the apparent value of things would change given one's immortality.

Recall my discussion earlier of the two types of individuals, the first who is in fact immortal but thinks they are mortal and the second who thinks they are immortal but in fact does not have immortality. Suppose I am living an immortal life but my personal experience is still essentially finite. Couldn't I still appreciate and create things as if I were mortal even if I, physically, were not, i.e. even if I in actual fact had external immortality? Again, Kass's worry seems to hinge on this notion of one's certain knowledge of one's immortality, if one could ever have certain knowledge that one will never cease to exist in the future. It seems reasonable to suggest that we can still live the physical life of an (contingent) immortal but still retain the psychological life and experience of a mortal life, and be able continually to appreciate and create things of beauty.

What the distinction between the two types of 'immortal' does highlight is the importance of the possibility of immortal personal time to any discussion of the desirability of immortality. If such a thing is possible then perhaps external immortality is only a prima facie problem to the extent that personal time goes along with it. Perhaps there is no real problem of external immortality at all given the possibility of endless personal time, which is the real target? I have argued in the Introduction that what we should be talking about when having a debate over immortality is a kind of contingent immortality and that both potential immortality of the organism or of the person have to always be considered under the blanket of contingent immortality. Thus maybe when it comes to the question of whether living forever would be desirable, it need only boil down to whether endless personal time would be desirable; that what actually matters is not that matter of fact of an actual infinite existence but only the personal experience of a potential endless life.

3.1.4 Loss of virtue

Finally, Kass argues that an immortal cannot be virtuous. Only an individual who is susceptible to decline and death can truly act in a moral way, or exhibit virtuous behaviour, such as courage, generosity, or devotion to justice, these values would evaporate if we knew too little of death.

It looks like Kass's argument is that, since it is only possible for the mortal to give their own life for what they believed was right, to be morally courageous, then the immortal is forever excused from this duty. Because Kass gives the argument as an argument against the individual who cannot die—i.e. has necessary immortality—there is a problem in counterarguing such claims, because even profound life extension will not mean invulnerability, which is what Kass seems to be thinking about at least in this argument. However, Kass's claim about the individual who knew less of death would be less noble or virtuous still has relevance in the context of life extension.

Is the inclination to be moral or show virtuous behaviour towards others contingent on being mortal and having a finite lifespan? Even those who had profoundly extended lives could give their life for what they believed in. but even if they were necessarily immortal and could not actually give their life in defence of what they believed was right, this does not necessarily mean that such an individual did not care about being moral or exhibiting other virtuous behaviour. Kass wants to claim that being mortal is a necessary condition for being virtuous but this does not seem to be the case.

There are other ways to live virtuously and exhibit virtuous behaviour other than the ability to give one's life if that is what was needed to protect what was right. In fact, other kinds of virtuous acts might become even more frequent given the greater opportunity to develop a virtuous character allowed by a much longer life. From an Aristotelian approach to virtue and moral excellence, for example, the virtuous individual practises over the years of his life the habit of acting virtuously and living well, and such an individual is not shifted easily from such moral excellence. Consequently, moral excellence is a life-long activity, in which the person acts in accordance with virtue over a lifetime. Given this, one might claim that the more years of life one had to practice acting virtuously and to develop a virtuous character the better that life was, for oneself and for others.

Kass wants to say that not only are longer lives bad for individuals, but that it is bad that individuals live longer lives. Individuals would suffer, which would be bad, but the fact that there were longer lives would also be bad. Kass's arguments have potentially demonstrated that it would be a bad thing for an individual to go through life thinking they are immortal or that they have unlimited personal time. As such they have not shown that life extension is a bad thing, i.e. that persons living longer lives is a bad thing. They have also not shown that it would be irrational for an individual to desire

profound personal survival because, among other things, such a life would be meaningless.

Kass makes a lot out of death's supposed connection to meaningfulness. But death's connection with a life's meaningfulness is clearly controversial. Let's grant a life without death is meaningless. Death, then, is the thing that gives meaning to a life only if that life is not exempt from death. But what if we are talking about the kind of immortal life that is not necessarily exempt from death, an individual who will potentially experience profound personal survival but is still possible for them to die. Is this individuals' life, then, not meaningless? If this person's life is not exempt from death, but simply has the potential to go on for a very long time, and death happening at some point in a life is necessary to give meaning to that life, then their life is meaningful. But this all seems a bit nonsensical. If the only reason that a lifespan which ran a thousand years is meaningful is because it ended in death, but a lifespan which ran a thousand years but was still ongoing, and for all we know is exempt from death, is not meaningful, then claims of meaningfulness based on death or limitation, specifically limitation in time, just seems arbitrary. Consider comments made by Robert Nozick about the dogged presumption of a connection between death and meaningfulness (1981: 580):

> …there is a general assumption that certain limits, certain pre-existing structures into which things can be poured, are necessary for meaningful organization…Even were this general assumption true, though, death constitutes only one kind of structural limitation: finiteness in time. Other kinds are possible too, and we well might welcome these others somewhat more. The dual assumption that some limitation is necessary for meaning, and limitation in time is the only one that can serve, is surely too ill established to convince anyone that mortality is good for him.

Along with the assumption that death is necessary for meaning Kass's claims about the benefits to us of our mortality also share a common idea: that the current human lifespan and limitations on human experience is the best possible state of affairs for human beings and that some other state of affairs would be less than optimal. His claims seem to imply that what we humans currently have is the ideal way to experience things within a human life, ideal for the human capacity to process and benefit from living the life of a human being. The introduction of life extension would change this, and not to the benefit of us. There is an analogue for this view in the psychological phenomena that occurs for individuals who adjust to their circumstance following a misfortune that has limited their capacities in some way. Think about those people who have suffered

an accident and have lost the use of their legs. They are naturally unhappy for a period of time immediately afterwards, but it is often the case that a survival mechanism of sorts kicks in and they adjust their preferences by their situation and their happiness levels return to a baseline and life without the use of their legs simply becomes the norm. We could reasonably suppose the same phenomenon in regard to the situation in which an individual is given the opportunity for a longer life. But these are contingencies I am talking about and I am far from suggesting anything universal, but they suggest at least that we might get used to living much longer lives and eventually consider it the normal state of affairs for human beings.

On this broader interpretation of Kass's claims he is implying that life extension would be a bad thing beyond simply being foolish for an individual to desire. But as I have suggested above, his arguments have not gone anywhere near far enough to demonstrate the badness of life extension.

3.2 Would profound personal survival be unattractive?

Bernard Williams (1973) has an even grimmer prognostication for the individual who desires profound personal survival. Not only would a radically longer life necessarily lead the individual to a state of permanent boredom if they remained largely the same kind of person throughout, but any attempt to avoid tedium by changing the kind of person one was—by altering the things one desired or modifying the personal goals one habitually pursued—would find the individual in a life that was of negligible concern or importance to the person that started out on that life. Either way, profound life extension should be an unattractive prospect for the individual, and not rationally worth pursuing.

Williams presents a model of profound life extension in the fictional case of a woman who has been given an elixir of life by her father when she was 42, preventing her from ageing and apparently making her immortal. This is the central character of the play 'The Makropulos Case', written by Karel Capek.[27] The woman, Elina Makropulos, though she has been known by different names throughout her long life, all with the initials 'EM', is now 342 and has done and experienced all the things that could

[27] First performed in 1922.

possibly be interesting and meaningful to her and finds the prospect of continued life unbearable.

Williams wants to argue that, though the example is fictional, it is not a peculiarity of EM that she found a profoundly longer life unbearably boring and deprived of meaning. Rather, he says, this would be a necessary condition of a potentially endless existence. As such "we could have no reason for living eternally a human life" (1973: 89). Williams tries to demonstrate this by presenting a dilemma for the budding immortal.

First of all, let me say that, although the target of Williams's argument is ostensibly 'immortality', if we alternatively take the case of EM to represent how things might play out for someone who uses life extension, then his target is to show that a life prolonged a few hundred years would be a bad thing for the individual. Of course, if he is right, then it follows that living forever is bad. But the stronger claim is that a mere few hundred years (like that experienced by Elina Makropulos) would necessarily be undesirable as well.

I will follow John Martin Fischer's (1994) useful framework of Williams's discussion of immortality. Williams says that two criteria must be satisfied in order for any model of immortality to be desirable to an individual: (1) immortality must preserve an individuals' identity throughout, and (2) immortality must be attractive to the individual throughout. Furthermore, within the second criterion Williams presents a dilemma. Either the individual's desires and goals remain the same throughout or they are allowed to change. If the individual's life goes the first way, they will eventually reach a state of unbearable boredom like EM; if the second way, they avoid tedium but it is unclear how the individual can decide whether this life is attractive to her or not because her familiar character—her way of judging the goodness or badness of her life—is so altered as to be removed from her.

Let's now bring this over to Williams's example of EM and her particular model of profound life extension, in order to see whether we can infer from her fictional case, as Williams does, necessary truths about the desirability of a profoundly longer life.

There are perhaps only a couple of points to be made about Williams's first criterion for an appealing immortal life and then it can be set aside. This is the criterion that the remote future person living the life must be sufficiently identical to me that started out on it. This is obviously a necessary requirement for a profoundly longer life to be attractive to me, because 'I' want to be the one living it in the future. Williams

thinks that during the course of a potentially endless life in which the individual is deeply engaged in activities where they can lose themselves and avoid constant reminder of the enormity of their indefinite existence, that they would become so removed from the person they once were that identity would have been violated without the individual realising it. Like an umbilical cut when one wasn't looking, casting the person that was adrift with little hope of reunification with the person that is, across a vast expanse of years. Others have also thought there would be a threat to the psychological grounds for identity if we lived significantly longer lives. Walter Glannon (2002), for example, argues that one's prudential concern for one's remote future self would be so weakened by a much longer lifespan, that you would have little reason to care about what happens to your future self and vice versa, any connection between mental states, desires, and goals of the original person tenuous at best.

In the case of Williams, what he is claiming is that the future person living the immortal life would be a different person to the one who began the life. Williams is perhaps making too strong a claim about the potential threat to identity. I would suggest that Williams is conflating the notions of 'changes to one's character' (e.g. one's desires, goals, values) and 'changes involved in having a succession of numerically different selves'. A person living an indefinitely long life by engaging fully with stimulating intellectual activities is perhaps negligibly different from the ordinary case in which a person's character, goals, and pursuits, change continually as she gets older and has new experiences, interests, and meets new people. The sameness of character is not a necessary condition for sameness of numerical identity.

The only way to make sense of Williams's views about personal identity in his argument presented in the Makropulos Case paper is to think of Williams as holding the view that significant changes to goals, desires, and character over a profoundly long life would necessarily result in the eventual non-existence of the person who started out on that life and the existence of a new person, in short, a change to numerical identity. For example, Williams says in the article (1973: 96):

> Some philosophers have pictured an eternal existence as occupied in something like intense intellectual enquiry...The activity is engrossing, self-justifying, affords, as it may appear, endless new perspectives, and by being engrossing enables one to lose oneself...But if one is totally and perpetually absorbed in such an activity, and loses oneself in it, then as those words suggest, we come back to the problem of satisfying the condition that it should be me who lives for ever...

Williams here is worried that, if one chooses to fend off boredom over the course of a profoundly longer life by engaging in activities (primarily intellectual activities) in which one can lose oneself (the only kind of activity that might sustain a perpetual existence, according to Williams) and constantly change one's preferences, interests, and character, personal identity will genuinely be threatened. It's not unreasonable, then, to infer that Williams's arguments presuppose a psychological criterion view of personal identity, and that he means to suggest (like in the passage quoted above) that, given enough time, engaging in novel and engrossing activates and interests will eventually result in a different person existing after, say, 300 years of life, from the person who started out on that life; that in the end it will not be 'me' who started out on the life. While Williams's view of personal identity deduced from his discussion in the Makropulos paper might not be Williams's all things considered view on personal identity over time, it is, however, the standard interpretation of Williams's discussion from this paper (see, once again, Fisher 1994) and it has become typical to interpret his view in the Makropulos Case article as implying that a different person, or at least a person sufficiently different from the one who started out on the life, will exist at the end of a life extended a few centuries, enough to make it irrational to extend one's life this long.

Williams seems to be saying that one would no longer be the same person in the numerical sense if one were to heavily involve oneself in such engrossing activities to stave off boredom over the course of an indefinitely long life. But this claim does not seem justified especially given the source of the threat to identity, which Williams says is "intense intellectual enquiry" activities that are "engrossing" and "totally absorbing" and that enable "one to lose oneself" (1973: 96). If I am 'lost' in an activity, for example, absorbed in an engrossing novel, I am self-absorbed temporarily but I can hardly say that I am no longer able to understand this activity as being part of my own experience and that it is now the experience of another. If I am interpreting Williams correctly on this point, he is of the opinion that identity is put in jeopardy rather effortlessly and is a constant threat to the individual given any substantial amount of time consciously, intellectually engaged, that we are the kinds of creatures that are very malleable in this respect, like an animal which constantly alters its fundamental make-up by mimicking its environment in order to avoid threats to its survival. But if this is the case then Williams is not talking about threat to numerical identity, i.e. persistence of the same object or thing or person, but merely threat to character, because even if

one's desires and goals radically change the person can be said to remain the same. Fischer thinks the distinction is one between the 'content' of the experiences one has, and their 'ownership', which Williams fails to see. It is one thing to say that you forget yourself by becoming unselfconscious when engaged in an activity or experience, but another thing to say that it is not 'me' having the experience, that I do not own the experience any more.

I would suggest that we should assume even a profoundly longer lifespan would not threaten identity in any significant way. And in any event, I have chosen so far in the thesis to lay aside problems of personal identity as outside of my analysis. I will assume that survival of identity will not be one of the problems faced by those who wish to prolong their lifespan, though there may be other problems for the individual.

Williams's second criterion, which asserts that an endless life must be attractive to the individual in so far as it cannot contain constant suffering or unending boredom or be unrecognisable as something the individual chose (that she is getting what she paid for), is more intriguing and the dilemma Williams's presents within it proves surprisingly troublesome.

Given survival, there are two ways a profoundly longer life could go. Either one's desires and character stay the same, or they change over time. Let's assume one's desires remain fixed throughout a life of a few centuries, like EM's did. The first horn of Williams's dilemma is that such a life will eventually run up against a state of permanent boredom. Furthermore, Williams argues that this state of permanent boredom would be a necessary condition of a life in which one's desires and character stayed the same, i.e. that tedium is guaranteed.

First we must understand why Williams believes that living too long will necessary lead to dissatisfaction. Williams's theory is that there are two kinds of desires that make up the lives of such creatures as us. Conditional desires and categorical desires. Most desires we have are conditional ones, desires for food, sensory stimulation, etc., in short, egoistical desires that one can satisfy fairly easily and that can only be satisfied on the condition that one is able to experience their satisfaction. Categorical desires are different and are not conditional on our being able to experience their satisfaction. Categorical desires are desires one has and wants fulfilled whether one is alive or not to see them to satisfied, and as such they that have the force to 'propel' one into the future, to make continued life meaningful, for example the desire to see one's child go to university. Williams argues that whatever categorical desires

one had when one chose to live an endless life would necessarily be exhausted if they remained largely fixed for the duration. This is what happened to EM three centuries after she drank the elixir, "everything that could happen and make sense to a particular human being of 42 had already happened to her…all the sorts of things that could make sense to one woman of a certain character; for EM has a certain character and…seems always to have been much the same sort of person" (1973: 90).

I will not scrutinise Williams's assertion that we actually have categorical desires, but assume this as being relatively uncontroversial given verbal reports we commonly get from people about the kinds of desires they have about their own lives. The question is, if one's categorical desires—the kind that make one's life meaningful and make continued life something to look forward to—remained largely consistent over the course of a few centuries, would they necessarily become exhausted, delivering a once meaningful life unto a state of permanent ennui?

Fischer believes that one could sustain one's desires and pleasure indefinitely. In much the same way as Williams, he suggests a distinction between two basic kinds of pleasures or activities that human individuals are capable of based on the nature of conscious experience. First we have 'self-exhausting pleasures' or pleasures that once satisfied one does not want or need to repeat, even at some remote point in the future, e.g. achieving some long term goal, such as finishing one's PhD thesis. Second, we also have 'repeatable pleasures', or pleasures that are fulfilling and satisfying at the time one experiences them, but that one also wants to experience again, repeatedly, at some point in the future, e.g. pleasures that come from eating good food, engaging in pleasant conversation, viewing a sunset, etc. There is a sense in which these repeatable pleasures might become boring if partook in too frequently rather than more sparsely, but they remain the kind of thing, argues Fischer, that might sustain indefinitely the life of an individual.

Fischer suggests that Williams, in positing the necessary boredom thesis of the first horn the dilemma, is concentrating too much or too exclusively on the self-exhausting kinds of pleasures, whereas if he considered the repeatable kinds of pleasures the prospect of an endless life doesn't look so unattractive. But Williams might respond that partaking only in the repeatable pleasures or desires, e.g. for food, entertainment, the viewing of sunsets, whose only value seems to be the immediate pleasure that we get out of them, are not the kind that give life meaning or that could necessarily sustain an indefinitely long life. The kind of desires that give life meaning

are those categorical desires that we want to satisfy for their own sake and that are related to us or the kind of person we are and want to sustain throughout life. When these kinds of desires run out then continued life is meaningless. Williams seems largely right about this. The question is whether Williams is right that categorical desires would necessarily become exhausted.

Wisnewski (2005) has suggested an interesting reply to Williams. He suggests that Williams has overlooked the possibility of renewable or open-ended categorical desires or of their re-emergence in the future. He uses the example of an immortal whose only categorical desire is to become the world's best musician. Many centuries later, the immortal has mastered every instrument in existence and it seems she has reached the end of this activity and the extent of this categorical desire. Thus, if Williams is right, perpetual boredom must ensue. But what if we expand the example, where, after a couple hundred years of boredom, a new instrument has been invented (not to mention the fact that music is conceivably endless, even with the limited number instruments the immortal had use of before), so the immortal's categorical desire has not actually been satisfied after all. Wisnewski's point is, firstly, that there are always going to be new things invented given enough time. Secondly, creativity and progress do not stand still, even if the immortal does. Moreover, the immortal cannot be certain that her categorical desires are really and truly exhausted because of the aforementioned possibilities. He says (2005: 34): "the example is meant to show that, even if one enters a state where no categorical desires are present, it does not follow from this that the state is permanent…Williams situates an infinite life in a finite set of activities". The utility of Wisnewski's example is to show that Williams is, in fact, making empirical assumptions when he argues that categorical desires will *necessarily* become exhausted.

Perhaps Williams is being mislead by his example of the isolated immortal EM, and is implicitly assuming that the immortal would be living in a static world with static people. He pictures the immortal consuming all the meaningful experiences she can until she has exhausted all of them and then must say "that that is all there is and will ever be for such a creature as me". Williams is perhaps unassumingly reaffirming the familiar Lucretian refrain, that the significant properties that make a human life good or meaningful are finite, and those that exist are always complete and no extension of these properties will somehow alchemically concoct new kinds. Or perhaps Williams means that there is some kind of fundamental baseline of human interest and desire that would not correspondingly be augmented by living longer, already prefiguring the end of any

individual who attempts to live indefinitely a life that must necessarily terminate in boredom.[28]

What then if our desires and character were allowed to change over time and so we avoided the famous state of permanent boredom? Would having different desires and character be a bad thing for the individual? I would suggest the real puzzle that the second horn of Williams's dilemma presents us is the question of whether you are made worse off or better off simply by having different desires and character as you live out a few centuries. Williams thinks he has shown that you are worse off, because they will not be 'your' desires and character. But there is an obvious sense in which the life of the remote future you with some kind of desires and character, will indeed be you or be attractive to you, simply by virtue of the fact that it is you we are talking about when we ask these questions. Your desires 400 years from now are 'your' desires, just a different selection of them out of the entire collection taken from your life as a whole; they are but the ones that merely instantiate themselves later on in your ongoing existence. Therefore, I say that simply having different desires over the course of a radically longer lifespan is, at worst, neutral, but it's hard to see that 'you' are made worse off.

Could you be said to be better off because life extension allowed you to have different desires within the same life? Perhaps you're not made better off either. Does the mere addition of extra desires to a life make you better off than you would have been, presuming we're talking about the kinds of categorical desires that give a life meaning? It's hard to see that it does. Such a thing is perhaps also of neutral value for the individual, your desires, at the time you are pursuing them, simply are the particular desires you have at that time. To have different ones at a different time, assuming they're desires that bring equal pleasure upon being satisfied to the previous ones you had, simply amounts to trying something different, neither better or worse, like desiring strawberry ice-cream when you once desired vanilla ice-cream. But having the *opportunity to* pursue different desires is, if anything, better than not having the opportunity.

[28] Bortolotti and Nagasawa, 2009, have drawn on empirical psychological data about situational boredom, a state of temporary boredom associated with situational factors, and habitual boredom, a chronic state of boredom correlated with an individual's character or personality traits, to suggest that an immortal life would not necessarily be afflicted by habitual boredom simply because of repeated experiences. Rather this more dire state of boredom, which corresponds to Williams's state of permanent boredom that he claims would necessarily befall every immortal individual, would likely be the result of traits contingent to certain individuals. Those individuals more likely to experience chronic boredom in a longer lifespan are those who find it difficult to chose life goals and be motivated by them and therefore would likely have reached a state of chronic boredom during the course of a regular lifespan.

Thus, contra Williams, you are not made worse simply by having different desires or character as you live out a few centuries of existence. So, of the two ways a profoundly longer life can go, this may be the way to go, since we avoid the state of permanent boredom, by allowing our desires and goals to change over time, which is no threat to our survival and does not make us worse off in any foreseeable way. Thus, life extension is at least not bad for you in terms of your interest in surviving individual life extension, and may, in the sense of expanded opportunity for trying on different desires, goals and values, be a good thing.

I remain undecided on the question of whether an individuals' capacity to sustain indefinitely categorical desires over an indefinitely long life. Part of me wants to suggest that this is largely an empirical matter to do with the nature of human psychology and its facility for satisfaction of preferences and the formulation of new ones, which is perhaps something that will change for human beings if we ever get the chance to test the temporal limits of desire.

However, this being said, I am satisfied that the alternative to living a life of fixed desires and character would not present a problem for the individual, certainly not so much as to make the prospect of this second kind of potentially immortal life unattractive. Allowing one's desires and goals to change might be the only kind of life possible for those wanting to live profoundly longer lives, and it is doubtful that Williams's claim that this kind of life would be in some way harmful, in the sense that allowing one's desires and character to change or expand over time would be unacceptable, should dissuade them looking forward to trying out a few centuries.

3.3 Conclusions

In this chapter I have examined arguments from two philosophers that profound life extension would be a bad thing for an individual and thus foolish for to desire. I found that the arguments of neither philosopher conclusively show that profound personal survival would necessarily be a bad thing. Thus an individual's desire for such survival is, at least from the arguments of Kass and Williams, fortified.

Even if Kass or Williams are right that it would be irrational for an individual to desire profound personal survival because such a life would ultimately be bad for them

(though I have concluded they are not), it does not follow that life extension is a bad thing. An answer to the question, 'Is the prospect of living a profoundly longer life a personally attractive one?' does not constitute an answer to the question, 'Are organisms with longer lifespans a bad thing?' or 'Should we develop life extension?' In this, the authors attempt to reach further than their arguments allow.

In the next chapter I will examine the question of whether personal survival is in fact a good thing for the individual.

CHAPTER 4

WOULD LIFE EXTENSION BE A GOOD THING?

> Death has stolen upon you unawares, before you are ready to retire from life's banquet filled and satisfied.
>
> Lucretius, *On The Nature of the Universe* bk3, 959.

I have argued contra Williams that the desire for a profoundly longer life is not necessarily irrational and contra Kass that our mortality is not necessarily good for us. I have claimed, on the other hand, that it is rationally permissible to want to live indefinitely and that a longer life is not necessarily bad for us. But this is only half of the story. There is a further question: is more life a good thing for a person, ceteris paribus, so that it is rationally permissible to want to continue living?

In this chapter I will try to give a positive answer for why wanting a longer life, in favourable circumstances, is rationally permissible and as a result, why prolonging one's lifespan might be a good thing. This will constitute one argument for why developing life extension—presumably by diverting funding and resources away from other human endeavours, technologies, or medicines—would be an important thing for humans to do and not merely gratuitous wish-fulfilment.

The most obvious way to proceed is to show why death, the thing that prevents continued life, is bad for the one who dies. If death weren't in some way bad for the person who dies, then using resources to develop life extending technology to prolong the human lifespan would probably be unjustified. The issue would quickly become merely one of resource allocation, where it would be less beneficial to invest in life extending technology if those resources could be redistributed to developing other technologies, say, for reducing more immediate human suffering.

I will not talk about the impersonal disvalue of death—how bad death is from an objective, impersonal point of view—but about death being something that is bad *for the person* who dies. I will also take death to mean our permanent non-existence. If death were not permanent non-existence and we survived in some form or other after physical death (a common metaphysical assumption of certain religions) then death

would not be a bad thing for anybody, because we would survive this kind of 'death'. But death is nonsurvival. And it is the fact that we do not survive our death that is why we generally think of death as a bad thing for the one who dies.

I began the last chapter by highlighting the tension between two strong intuitions that we have, (a) that death can be a good thing because living too long might eventually become disagreeable, and (b) that death can be a bad thing because it deprives us of continued life. We usually think that death is a bad thing for the one who dies if it comes before one has had the opportunity to do and to experience all the things one wanted to do and experience. It is plausible, then, to assume at least this much: that some people are motivated to prolong their lifespan from the desire to avoid premature death. The idea of avoiding a premature death points us towards one explanation of why it is good to have more life rather than less, because we are capable of viewing our lives, and of having desires about our own lives, in such a way that it is unified, so that more life allows for, among other things, the satisfaction of a person's desires, goals, and unfinished projects.

But not everyone has thought that death is a bad thing for the one who dies, or that having more life is any better than having less life. In this chapter I will begin by examining the famous Epicurean argument that death is not a harm to the one who dies, and follow this up by examining a less well-known Epicurean argument that suggests that a longer life is no better than a shorter one. I will explore different ways of responding to the Epicurean argument about the non-badness of death and ultimately conclude that although Epicurus is right to suppose that it does not make sense to say that the state of being dead is a harm to the dead person (and maybe this is all Epicurus wants to say) this alone is not sufficient to show that the event of death is not a bad thing for the one who dies. I will suggest, exploiting views derived from Nagel and Williams, that the harm of death is the event of a person's death—death thought of as that moment in which a person goes from existing to not existing and are deprived of the opportunity for more life—and thus that death is a bad thing for the one who dies because of the desirability of the good things that are taken away or frustrated by the event of their death. Epicurus is wrong because he fails show that the event of death is bad for the one who dies.

In the argument examined in the second half of the chapter, Epicurus urges us to think of ourselves as mere bundles of sensations, as capable of being complete and satisfied at every moment, at least in terms of our experience (pleasurable sensations

being the only intrinsically goods things anyway). However, it is possible for a person to view their own life not just as a collection of momentary experiences or sensations that happen to a passive subject, but as a unified whole, as the life of an active agent in the world able to reflect on the goodness or badness of their life. This is a perspective from which the person can view the event of their death as a bad thing because it is an event that they might not want to happen, an event that prematurely cuts short preferences they might have had about their future, goods they might have enjoyed had their existence been extended. What is perhaps doing the work is the particular sense of self that you have, an episodic sense of self or a narrative sense of self. A further question then might be posed: should we have an episodic sense of self if reductionist metaphysical view of the self is true? Perhaps one cannot help what sense of self they have, so there is no epistemic obligation on one to change. Answering this question will have to be left for another occasion, but is something that can be left in the back of the readers mind as they think about the issues addressed in the second half of the current chapter.

4.1 Is death bad for the one who dies?

In his *Letter to Menoeceus* Epicurus urges the reader to understand that their death will not harm them because they will no longer exist to be harmed by it, consequently we should think of death as "nothing to us" (Saunders 1966: 50):

> Accustom thyself to believe that death is nothing to us, for good and evil imply sentience, and death is the privation of all sentience…Death, therefore, the most awful of evils, is nothing to us, seeing that, when we are, death is not come, and, when death is come, we are not. It is nothing, then, either to the living or the dead, for with the living it is not and the dead exist no longer.

It might look to many readers like Epicurus is making a trivial point: Your death cannot harm you because when you are alive you are not dead and when you are dead you don't exist. Let's call this the 'trivial conclusion'. While we might readily accept the trivial conclusion, the Epicurean argument seems to want to say that death is in no way a bad thing for the one who dies. This we are not so ready to accept. Most of us think that death is frequently a bad thing for the one who dies, for example, if the one who dies is very young.

Are we merely confused? Is death is in no way bad for the one who dies? Or is death bad for the person who dies in some other way not identified by Epicurus's argument? This is not to say that Epicurus was unaware of the common feelings about how death can be a bad thing for the one who dies, for example, that it is a misfortune that the person who dies young misses out on a longer life or that anticipations and other psychological features influence one's feeling about death being a bad thing. On the contrary, Epicurus's motivation for the argument about the non-badness of death was to alleviate the natural fear we people have of dying.[29] But Epicurus thought that these ordinary views about death's badness were mistaken, in particular the worry that one will be directly harmed by one's own death.

Epicurus's argument might only seem to be clearing up a basic confusion that some people have that they will somehow be alive to experience some direct harm from their own death. But most commentators have thought that in order to say that death is bad for the person who dies in any way we have to directly refute Epicurus's argument. I will present multiple interpretations of the Epicurean argument throughout the chapter and address each of them in turn, building towards the conclusion that death can be a bad thing for the one who dies and that the various Epicurean arguments are unsuccessful in showing that this is not the case.

Epicurus's argument derived from the passage quoted above can be interpreted in different ways. I will define first what I mean by various terms that are used by Epicurus and will be used throughout the chapter. When Epicurus talks of death as something that 'when we are, death is not come, and, when death is come, we are not', I would initially suggest that he is talking about the 'state of being dead', not the 'process of dying', nor the 'moment' of death, if it is, in fact, a moment. That is, 'the period of non-existence after one ceases to exist'. If this interpretation is correct then what Epicurus argues is that the state of being dead is no harm to the dead person because it will involve no harmful experiences for her. To avoid confusion later on I will use the term 'dying' as referring to the process of decline of an individual. The term 'death',

[29] The motivation behind the Epicurean view can be given an historical explanation. Epicurean doctrines such as the argument about death were offered largely for 'therapeutic' benefit rather than any attempt at the time to argue that a shorter life is actually just as good as a longer one. In the Letter to Menoeceus he says: "therefore a correct understanding that death is nothing to us makes the mortality of life enjoyable". One ought to practise thinking about life in this way, in order to lessen strong feeling that one's death is a bad thing. With the strongly hedonistic nature of his doctrines Epicurus had the goal of setting out a comprehensive guide for a life free from needless worry about mortality. Contextualised as such, Epicurus's aim was one of helping people to enjoy whatever quality or duration of mortal life they do have.

because of its ambiguity in the English language, in which it could mean 'the state of being dead' or 'the event of death', I will clarify at certain points in the discussion when I am referring to death as 'the state of being dead' or 'the event of death' and how this changes how we interpret and respond to Epicurus.

To put an even finer point on Epicurus's argument I would like to suggest that it does not rule out, (i) that one's death is bad for others, which it can be, for example, when it involves the loss of someone close to you and the accompanying grief or psychological harm, (ii) that the process of a person's dying can be bad for them, e.g. a painful experience, which it obviously can be, nor (iii) that one's being dead is bad in a way that something can be bad but not for anyone, e.g. the impersonal badness of the fact that there is less wellbeing in the world than there could be.

The Epicurean argument, based on the first part of the quotation above, "death is nothing to us, for good and evil imply sentience, and death is the privation of all sentience", and given the interpretation of the term death above, can be formulated as follows.

The First Epicurean Argument:

(1) A state can be good or bad for me only if I experience it.
(2) Death is not a state I experience.
(3) Therefore, death (the state) cannot be bad for me.

The first Epicurean argument is this: the only things that can be good or bad for a person are good or bad sensations (I will refer to this as the 'experience requirement' for the harm of death); death (the state) is not a good or bad sensation; therefore death (the state) is not bad for the person who dies. Premise (1) derives its plausibility from an attractive view of wellbeing, i.e. hedonism, the view that the only intrinsic goods are pleasurable sensations and the only intrinsic harms are painful sensations. Thus Epicurus thinks that if we are convinced by this argument, that the only harms are painful experiences and we cannot possibly experience our death, it should be a great weight off our minds: our death cannot harm us.

However, most of us retain the belief that death is bad for the one who dies, for the reason that my death takes away all possible opportunity for more life. It seems here like we want to say that being dead is bad by comparison to some other state, the state

of being alive. If it makes sense to say this, then not all goods or evils that can befall a person are good or bad sensations. So premise (1) above is in tension with a commonsense belief view that something can be harmful to a person though they don't experience it. Epicurus believes only in harmful sensations, but we seem to believe there are such things as non-experiential harms that can happen to a person as well.

Thomas Nagel (1979) is suspicious that anything that can be good or bad for a person has to be good or bad experiences, and suggests a number of counterexamples involving states that can be good or bad for someone though they involve no actual harmful experiences for her. For instance: we can imagine a situation in which a person is ridiculed behind their back; or perhaps she is betrayed without her knowledge; or she loses all of her higher mental functioning following a brain injury and is reduced to a state of something like a contented child.[30] Any person would consider these things bad for the person were they to happen even though never involve any directly harmful sensations for the person. These seem like knockdown counterexamples to the first Epicurean argument, which stands or falls on the central claim that the only things that can be a harm to a person are harmful experiences. But it seems like there are harms that do not depend on the person experiencing the harm.

We tend to believe that the above counterexamples of harmful states or events are bad for the person because if they could somehow see the future and know what would befall them then they would desperately want to avoid them. Why must think that any possible harm that can happen to a person must occupy the same space and time as the person? Nagel's examples of non-sensational harms are relational or comparative harms that you can be the subject of but do not necessarily have to exist at the time of the harms themselves.

The 'experience requirement' premise of the argument might also not be as well motivated as Epicurus counted on, because even the hedonist can disagree with premise (1). Hedonism doesn't necessarily imply the experience requirement of harm. Think about the example in which a hedonist is rendered unconscious for an hour. During that hour she could have been having pleasurable experiences, but misses out on having these sensations because she is unconscious. It seems like the hedonist has been harmed by being rendered unconscious, because having pleasurable experiences is intrinsically

[30] Aristotle, in the *Nicomachean Ethics*, also talks about the kinds of contingencies that can affect one's happiness and makes a similar argument that it is thought that good and evil things can happen to a dead person "since they can happen to a person who is alive but not aware of them", for example, dishonours done to a dead man's children or descendants (Book 1, 1100a).

valuable for her and she is deprived the opportunity for having pleasurable experiences when she is unconscious.

But Epicurus is not down and out yet. He has a second argument. He says: "death…is nothing to us, seeing that, when we are, death is not come, and, when death is come, we are not" which suggests the following argument.

The Second Epicurean Argument:

(1) If P is dead then P does not exist
(2) If P does not exist then P is not in a harmed state
(3) Therefore, if P is dead P is not in a harmed state
*(4) Therefore, death is not a harm

In the argument made up of the first three premises, Epicurus is saying that death cannot be a harm for the one who dies because one doesn't exist in a state of being dead. Remember that I earlier called this the 'trivial conclusion', and one which we can accept. But Epicurus is making a further inference: that death is not a harm. This further inference I will disagree with.

Epicurus's second argument is stronger than his first because it doesn't appeal to what I've been calling the 'experience requirement' for death's harm, but endorses what can be called the 'existence requirement' for harm instead, which is expressed as premise (2) of the second argument. A person has to exist in order be harmed. While Nagel's examples of harms that are not experienced by the person are sufficient counterexamples to Epicurus's first argument, that is, counterexamples to the truth of the experience requirement premise, they are nonetheless all examples that involve the subject of the harm still existing. We might make an exception for the brain damage case, in which it could be said that the original person has ceased to exist following the injury and is replaced by the 'contented child'. But there is still a subject in existence who is perhaps harmed in the example.

The existence requirement, on the other hand, looks like a logical truth, one that resists appeals to any theory of wellbeing one might marshal against it. It is perhaps simply true that if a person does not exist they cannot be in a harmed state. Although that being said, later, when I discuss Williams's contribution to the response to Epicurus, his stressing of the importance of categorical desires, we will see that perhaps

Williams doesn't think that the existence requirement of death's harm is a logical truth. This is because, under Williams's description of categorical desires, your categorical desires in a sense 'live on' after your death, and thus can still harm the person in a comparative way even in a state of nonexistence, i.e. being dead. But I will not rely on this particular counterintuitive aspect of Williams's interpretation of categorical desires. I agree with the existence requirement of death's harm but argue that the harm of death comes earlier.

So we can concede to Epicurus that he is right that it is not bad for the person to be in a state of being dead. However, I want to argue that, while this is true, it is still bad to *die*. I want to claim that death is an event in a person's life. The 'event of death' is that moment which takes one into the state of being dead and it is this event that harms the one who dies. I have argued up until this point that what death refers to is the state of being dead. But this is not all that death is; it is not the death we are talking about when we say that death harms the one who dies. 'Death', the thing that harms, is an event that takes a person from existing to not existing. When the person is dead, death has already happened. So death is an event that happens to a living person in that final moment before they are in a state of non-existence. What we mean by death, then, the thing we want to say "harms" someone is an ante-mortem event, the very last ante-mortem event possible, because the subsequent state is being dead, with nothing in between.

I say that "death is an ante-mortem event", which the reader might naturally interpret as saying that death is something that occurs before death. But how can something happen before it happens? The statement is not as paradoxical as it might seem at first glance. The view that I defend in this chapter is that 'death' is ambiguous between a 'state' and an 'event'. When asking whether death is a harm, most philosophers have in mind death as a state—the state of being dead—and this assumption leads to many familiar paradoxes. But 'death' might also refer to the event that transforms a living person into a dead person. Thus, I claim that death is something that can only happen to a living person and cannot possibly happen to a person who is already dead. Interpreting death in this way might avoid certain paradoxes related to how bad death is. In any case, interpreting death as the event that takes a living person into non-existence seems to have been overlooked by commentators in the literature on the 'badness of death', including Epicurus, who can be reasonably interpreted as using 'death' as referring to 'the state of being dead', which I claim is a state after the 'event

of death'. If death is interpreted as the state of being dead, then it is truly nothing to us. However, I have tried to show in this chapter why death seems (in most cases) to be a harm, and argue that it is a harm because death is an event in the life of the living person, thus an event that can be a harm to them. When a person is no longer capable of experience, then death has already occurred. So, to reiterate, one's death is an event that occurs before one is in a state of being dead—thus, death is an ante-mortem event. Death (the event) harms the one who dies because it is the thing that deprives them of continued life.

So the living person is harmed in that final moment before they die by having any further opportunity for good experiences, preference satisfaction, etc., deprived of them. Thus it is bad to die, or put another way, the 'event' of death is bad for the one who dies. Therefore, in response to Epicurus's further inference, death is a harm for the one who dies: they are harmed in that final moment of existence in which everything is taken away from the person. This is the event of their death, and death as an event in a person's life is what I will mean by death from now on.

Nagel uses his examples of how a person may be harmed without experiencing harm in support of what is now known as the 'deprivation account' of death's badness. On this view death is bad not because there is anything intrinsically bad about being dead but bad because of the desirability of what death deprives the person, i.e. the goods of life, or at least the opportunity for further goods. I agree that the badness of death (the event) for the person comes from when they are deprived of the goods of life, in whatever way one interprets those goods based on their particular theory of wellbeing, e.g. the opportunity for more pleasurable sensations, further preference satisfaction, etc. But Nagel's mistake is to try to show that the harm of death by deprivation of goods is somehow a comparative harm: that a person, in the state of being dead, is worse off (and thus harmed) by being dead than they would have been if they were in a state of existence. This approach gets Nagel, and others in favour of his version of the deprivation view, in trouble when challenged to show how the person can be harmed in a state of nonexistence. This is exactly the challenge presented by the core part of Epicurus's second argument, or the existence requirement of death's badness.

I am sympathetic to a deprivation view of why death is bad for the one who dies, but I accept the truth of the existence requirement for death to be a harm. My view is a deprivation view but it manages to resolve a couple of problems Nagel faces with his view. Firstly, Nagel's view is that being dead is bad for a person in so far as were they

alive they would have been enjoying the goods of life, so they are harmed in a comparative way even in a state of nonexistence. Thus he must reject the existence requirement, a move that is very counterintuitive. I don't need to deny the existence requirement, because on my view, death is an event in the person's life, the final depriving event of their life, in which the person is harmed by having the goods of life (as well as the opportunity for more goods of life) taken away from them. This depriving event is an event that happens to an *existing* person. Secondly, Nagel is forced to say that there is no time at which death is a harm for the one who dies, that the badness of death for the one who dies is 'timeless'. On my view I am able to point to a time at which death is bad for the one who dies, that time is the event of their death, the depriving event, which I am able to say happens to an existing person because it is that final moment of life when they are deprived of more moments of life (and the chance for more sensation, or preference satisfied, etc.). This is the time at which death is a harm to the one who dies.

Consequently, my view can be called *the deprivation event view* of the badness of death, while Nagel's can be referred to as *the deprivation state view* of the badness of being dead. The deprivation event view is able to show how death is bad for the one who dies, while avoiding the problems faced with the deprivation state view. Nagel's mistake is focusing on the state of being dead rather than the event of death or the moment that the person dies in which they go from existing to non-existing. This is more precisely the point at which the deprivation, the harm of death, happens to the one who dies.

Epicurus says that death can only be a harm when you're alive or when you're dead, but it is not a harm in either. But Epicurus is wrong about the first horn of this claim, because death is a harm when you're alive. Can Epicurus answer this new challenge? What if we reformulate the Epicurean argument so that death is no longer 'the state of being dead' but the 'event of death' or the moment the person 'dies'.

The Third Epicurean Argument:

(1) If P dies then P does not exist.
(2) If P does not exist then P does not suffer harm.
(3) Therefore, if P dies then P does not suffer harm.

But I do exist during the event of my death, which is a last moment of existence, the moment I die. I am harmed in this last moment of existence by being denied the opportunity for existing for more moments, for being denied a future. Therefore, premise (1) of the argument above is false. The subject being harmed is a living person; it is that living person who a moment ago had a life ahead of her, which has now been taken away.

Both Nagel's deprivation state view and my deprivation event view might be seen to have the further implication that a person's death is worse the more life is deprived the person when they die, other things equal. Not only is the person harmed by dying, some deaths are more harmful than others because the event of death happens sooner rather than later, and thus they are deprived of more life. This aligns with our general intuition about the relative badness of the premature death that befalls say a young person and what is thought of as the less premature death of older persons.

But it's not clear that either view exactly implies that some deaths are worse than other deaths depending on when the victim of death would otherwise have died. What does the event of death take away from you: the benefits of life you would have enjoyed had you not died. But on my deprivation event view this doesn't mean that the person's being dead continues to be a bad thing for them after the event of death. However the view might allow that the harm of the event of death (which last only for a moment) can be of lesser or greater magnitude to the person who dies, lesser harm to the person who is satisfied with the life and opportunity to enjoy the benefits of life that they've had up till the point of their death[31], greater harm to the person who dies well before they have had the opportunity to experienced the goods of life or satisfy certain long-term desires, and so on. I will not commit the deprivation event view to whether or not it implies that the event of a person's death can be of greater or lesser harm to the one who dies, I only want to show in the present chapter that the event of death *is* a harm to the one who dies. I will address the issue of whether the death of a younger person is worse than the death of an older person in Chapter 8 when I examine the notion of a 'fair innings' and whether older persons have just as much claim to continued life as younger persons.

[31] The idea of a premature death is not necessarily guided by the number of years lived by the person at the time of death, but could also be influenced by the amount of opportunity to experience the goods of life, or the capacity to enjoy the benefits of living, etc., that the individual has been allowed.

Nagel's deprivation state view, on the other hand, might imply that all deaths are premature, that no matter when the person is deprived of life they could have had an indefinite amount of life ahead of them. Nagel in fact believed that death is a premature termination of indefinitely ongoing and continuous goods and preferences and projects. He says (1979: 9-10):

> Observed from without, human beings obviously have a natural life span…A man's sense of his own experience, on the other hand, does not embody this idea of a natural limit. His existence defines for him an essentially open-ended possible future…Viewed in this way, death, no matter how inevitable, is an abrupt cancellation of indefinitely extensive possible goods.

So Nagel's view is that maybe our life is never complete from our own perspective of it and every death is a premature end to possible ongoing experience. This contributes to Nagel's deprivation view that a person's 'being dead' continues to be bad for them after they die, and perhaps never ceases to be bad for them so long as they no longer exist when they could have otherwise existed. My deprivation event view however says that the person is harmed by only that event of deprivation when they die but not that the state of being dead continues to be bad for them after this initial (and final) bit of deprivation harm has been inflicted on them.

A supplementary theory of wellbeing might help to support the deprivation view of death's badness for the one who dies, and perhaps shed some more light on the intuition that some deaths are worse than other for the one who dies. The deprivation account points us in the right direction as to why death is bad for the one who dies, because death deprives the person of goods that were desirable. But it leaves open the question of motivation: why it is that we want to retain the goods that death takes away. One answer comes from a theory of wellbeing as the satisfaction of preferences; in particular, we have a strong reason to continue living in order to have our categorical preferences satisfied. Bernard Williams has argued that a person's death can be bad for them they had a categorical desire for something to happen or to obtain in the future which death frustrates. This is one reason why we can say that a person's death was premature and thus a harm to them.

Recall from Chapter 3 that Williams makes a distinction between conditional desires or preferences and categorical desires. It is true that many of the desires we have are contingent on our being alive, e.g. the desire for food, engrossing conversation, sensory stimulation, etc., things I want so long as I'm here. In short, they are desires

whose satisfaction and importance are limited to or conditional on our being alive. However, as reflective agents living a life, thinking about our future and having preferences for how we would like it to be, we seem to have another kind of desire, usually long term preferences that we organise and plan our lives around and which give us a reason for wanting to go on living, e.g. the desire that one's children go to university, or that one's long-laboured-over novel gets published, etc., they are desires that, as Williams puts it, 'propel' us on into the future and lend meaning to our lives as a whole.

Williams's preference theory of wellbeing assumes a view of wellbeing that is more sophisticated than the Epicurean view, which is a general hedonistic view, that the only intrinsic goods are pleasurable sensations. Williams's view presupposes that the satisfaction of desires is the most important good that a reflective agent can enjoy or that can be deprived of them. Having one's desires satisfied is ultimately the good of continued life or continued consciousness for the person.

The Epicurean view cannot explain why the event of death affects one for the worse. According to my account of death the person is harmed by the event of their death and on Williams's preference satisfaction view of wellbeing, the event of death harms the person because it frustrates the categorical desires they might have had. It is not that death cannot be bad for the hedonist, we have seen that it can, because the event of their death deprives them of the opportunity for more pleasurable experiences. But for those with categorical desires, not only conditional ones, the event of their death is more acutely bad for them because they had a stronger reason to want to continue living than the hedonist. Or perhaps we can put it this way, those with categorical desires had already staked their claim to more life. They had already projected themselves into the future, a future in which they would be living their particular life satisfying their particular desires. While those with only conditional desires need not reflect on a world where they cease to exist, for them, continued life is only desirable at all at a time when they are alive, but as soon as they ceased to exist 'all bets are off'.

Thus the Epicurean conclusion about death, that it is in no way bad for the one who dies, seems unsatisfactory for most of us because we are able to reflect on the badness of having one's categorical desires frustrated by the event of death. This is why death is a misfortune for the person who dies and can be said to be 'premature' so long as they had unfulfilled categorical desires when they died.

Williams claims that an important aspect of categorical desires is that they do not depend on our being alive to be satisfied. Williams might want to say that any categorical desires the person had before they died 'survive' in a sense after their death in how they wanted their life to be had they lived. This 'residue' of badness that lingers after the death of an agent who had desires almost allows the person and their death to coexist historically, in the sense that we can say 'this is how the person's life would have gone had they lived and how the fact that she didn't live is a bad thing for her'. However, I will not hold with the counterintuitive implications of Williams's particular description of categorical desires, i.e. that they still harm the person even in a state of being dead. The Epicurean could challenge Williams with the existence requirement: how can it be bad for the person who no longer exists to have their categorical desires frustrated? In the absence of a subject of existence there is no person to be harmed. And since any categorical desires one might have had cease to exist along with them, as well as any harm that might have come from seeing one's desires frustrated. It is enough to agree, along with Williams, that the frustration, the cancellation of these future-orientated preferences the person had by the event of their death was a harm to them at the moment of death.

Combining my version of the deprivation account and Williams's suggestion of the importance of categorical desires provides a response to Epicurus's challenge that death is not bad for the one who dies, by showing that the harm of death is at the event of the person's death because in this final moment of existence death (a) deprives the person of a longer life, and this is further emphasised by (b) that death frustrates any categorical desires they had. Death can certainly be bad for the one who dies. Thus the Epicurean view of death is revealed to be false because it fails to recognise the possibility that a person's death can harm them and in what ways it is a harm. All of the Epicurean arguments presented above fail to demonstrate that death is not a harm to the one who dies. Therefore, they must be rejected.

A person's death can be a bad thing for them. But neither my deprivation view nor Williams's view commits us to the conclusion that a person's death is always bad for them. Being deprived of the goods of life assumes that the life would have had goods. Such is not always the case, for example, in extreme cases of a life a constant torture, or some other kind of suffering. In such examples death might be a blessing for the person and the only thing it would be depriving them of are painful experiences. And Williams argues (see Chapter 3) that death can sometimes be a good thing because

it can be a good thing not to live too long if one's categorical preferences have all been exhausted, making continued life meaningless.

Epicurus is right to conclude that a person has no harmful experiences when they are dead, but this establishes nothing that any rational person would disagree with and fails to get at the heart of why we generally think a person's death is bad for them. Most of us are not merely subjects of experience but are active agents in the world living a life and able to reflect on the goodness or badness of that life. This makes it possible that harms that can befall a person are wider than mere unpleasant sensations because persons, with a certain view of their self, have the capacity to view their prospect death if it will be premature as a bad thing for them. So death can be a bad thing for the person in a way that it is not for a mere passive subject of experience.

4.2 Is a longer life better than a shorter life?

Death is bad for the person who dies not simply because of the good thing it ends (life) but because of the good thing it prevents (more life). More life is better for a person than less life, other things equal. Epicurus, however, had a challenge to this claim too. Not only did Epicurus argue that a person's death is not bad for them but he also concluded that the desire for more life is unnecessary, it is, in fact, redundant so long as one recognised that one can get sufficient pleasure from a short time alive as a long time. This additional Epicurean argument can be thought of as a further response to the deprivation/frustration of categorical desires view of death's badness and its implication that so long as being deprived of more life is bad for the person, then more life is better for the person than less life. We find Epicurus's views about longer life in his 'Principal Doctrines' (Saunders 1966: 55):

> Unlimited time and limited time afford an equal amount of pleasure, if we measure the limits of that pleasure by reason…the mind, grasping in thought what the end and limit of flesh is, and banishing the terrors of futurity, procures a complete and perfect life, and has no longer any need of unlimited time. He who understands the limits of life knows how easy it is to procure enough to remove the pain of want and make the whole of life complete and perfect.

Given the equivocal nature of the passage above, Epicurus's view might be brought to light by presenting the opposing argument. Take L1 to refer to 'Life 1' and L2 'Life 2':

(1) If L1 is longer than L2 then, all else being equal, L1 contains more pleasure than L2.

(2) If L1 contains more pleasure than L2 then it is better to live L1 than L2.

(3) If L1 is longer than L2 then it is better to live L1 than L2.

Epicurus is challenging either premise (1) or (2) based on how you interpret him in the passage from the Principal Doctrines. Let's begin by interpreting Epicurus in the Principal Doctrines as challenging premise (1). When Epicurus says, "unlimited time and limited time afford an equal amount of pleasure, if we measure the limits of that pleasure by reason", he seems to be making the claim that the same amount of pleasure is only possible regardless of the amount of time the person has lived for. That is, if I had but a moment, then this would be a complete and satisfying life, as complete and satisfying as the life of the hedonist who has lived for a hundred years. This, on the face of it, seems clearly false.

But perhaps Epicurus is only claiming that we would do best to *think* of our experience as complete and satisfying at every moment, so as not to desire a life of longer duration, when he says "the mind, grasping in thought what the end and limit of flesh is, and banishing the terrors of futurity, procures a complete and perfect life, and has no longer any need of unlimited time". Epicurus might be seen as challenging premise (2) and arguing that it is not better to live the longer life because we can practice a way of thinking, of having a particular sense of self experience, that finds all the pleasure it needs given limited time, even if it is impersonally true that the longer life contains a greater amount of pleasure than the shorter life.

I think we can interpret Epicurus as claiming the second view: that we should practice having a sense of self that sees pleasure, or experiences of any kind, made greater by duration as an illusion, making the desire for more unnecessary. The Epicurean committing herself to this doctrine should find no additional or greater pleasure to be had outside the confines of a moment, and in this way, the Epicurean can 'banish the terrors of futurity' and be able to 'procure a complete and perfect life' no longer slave to desire for unlimited time, that it is not better to life the longer life than the shorter one.

Epicurus is likely not arguing for a metaphysical conclusion but rather giving a normative argument for how we should experience our lives, how it would be better *for*

us to experience our lives, based on a particular view of self-experience I will refer to as 'episodic'. This can be recognised as a rather unorthodox sense of self-experience implicit in Epicurean philosophy, but as we will see later, it is a view that has its contemporary supporters, and is in contrast to those who advocate a 'diachronic' or 'narrative' view of self-experience. But what is the episodic view of self-experience and is it better for a person to adopt than some other view?

The Epicurean view of self-experience is similar to what has been called having an 'episodic' conception of self-experience. Galen Strawson (2004) argues in favour of this kind of self-experience against 'diachronic' or 'narrative' conception of one's self, which is to think of our self as an agent in the world who has a whole history, who can project their desires into the future. But we can think of our selves and our experience in a very immediate episodic sense, without past or future or at least not connected to 'this person's' past or future in any meaningful way. Strawson articulates the episodic view in the following passage (2004: 433-34):

> …it seems clear to me, when I am experiencing or apprehending myself as a self, that the remoter past or future in question is not my past or future, although it is certainly the past or future of GS the human being…As for my practical concern for my future…it is biologically – viscerally – grounded and autonomous in such a way that I can experience it as something immediately felt even though I have no significant sense that [I] will be there in the future.

Similar to the Epicurean sense of self presented in the passage quoted from the Principal Doctrines, the episodic self, or the 'I' at this moment, is not particularly linked to this person's past or future events, nor does 'the 'I' at this moment' find much value or motivation in events outside moment-to-moment experience. If the experience of pleasure, for example, has a 'static' value rather than a 'time-dependent' or 'duration' value, then one's experience might always be seen as complete from this perspective. If we experience pleasure only contingently connected to temporality, then from such a perspective of self-experience more duration does not enlarge pleasure. There is no value outside moment-to-moment existence, and as such, coveting a longer life is simply pointless.

We can see that the episodic experiences life similar to a reductionist metaphysical view of the self akin to the Humean view of our selves as merely 'bundles' of experiences is true. Putting this into practise as a psychological view—that there is no unified self or agent and we are merely a collection of loosely connected

experiences, sensations, and mental states—a person might go about their life something like Strawson's description above of his subjective experience, moderating their desires so that they were purely conditional on their moment to moment existence. All worry about one's imminent demise vanishes if one has not projected oneself into the future beyond the immediate present.

But is this the way most of us go around thinking about our lives and having preferences that involve our future? The episodic way of experiencing life seems at odds with the basic nature of human desire and any idea of agenthood and of living a life connected in meaningful ways to 'my' past and 'my' future. I would suggest rather that we make sense of our life and our eventual death from an anti-Epicurean or anti-episodic view our self-experience and our lives. Furthermore, from the anti-episodic view of self-experience, we are able to see how death can be a bad thing for the one who dies and why continued life is valuable.

The value of a person's life and of its continuation comes from generally narrative view of their lives which connects what is probably tenuously connected sensations and mental states into a unified life of an agent, who is capable of thinking about their own future and having preferences that it be some way rather than another. Most of us are not mere passive subjects of experience, but active agents living a life that deprivation of which can be a bad thing for us. Wanting to postpone one's own death because one has not yet completed all one's projects and satisfied all one's desires, participated and enjoyed the relationships one wished to enjoy, can be seen as a rational aim perhaps only on an anti-episodic sense of self. Clearly, then, a non-episodic, (typically) narrative view of one's self, plays a vital role for many people in their evaluation of their lives and their decision of whether to extend their lives (Bunn 2009).

So what can we say about Epicurus's challenge presented in the passages from the Principal Doctrines? If we are anti-episodics we must reject Epicurus's challenge and accept premises (1) and (2) and conclusion of the argument above because having a longer life can be experienced as having the potential for greater pleasure than a shorter life. For an anti-episodic, pleasure is allowed to extend, to increase or decrease, desires able to be fulfilled or thwarted, because the person or agent themselves is extended over a period of time, that is, not only exists but has experiences that extend for longer or shorter periods of time. The episodic person and their desires, on the other hand, stretch perhaps no longer than a moment.

It is difficult to see how the episodic psychological view could be sustained, moreover, how it could be beneficial to a person. If more life or more experience is better for a person than less life or less experience, which I think we have shown in our conclusions about the badness of death, then the episodic view, which I've claimed is the same or similar enough to the Epicurean view interpreted from the Principal Doctrines, must be rejected. Or perhaps we might give a more moderate conclusion: death is a bad thing for the one who dies and more life is good, so long as one is not an episodic. Whether it is good or bad to be an episodic can be set aside, though Epicurus thinks it is a good thing to be an episodic so that death isn't bad for the one who dies, but the episodic must at the same time deny that having more life is good. Thomas Ligotti (2010: 41), for instance, thinks that we would be less beleaguered if we just gave up the notion that we are persons leading a 'life' and simply saw ourselves as conscious subjects who have an 'existence':

> Perhaps we may gain some perspective on our earthly term if we stopped thinking of ourselves as beings who enact a "life"…Instead, we should substitute "existence" for "life" and forget about how well or badly we enact it. None of us "has a life" in the narrative-biological way we think of these words. What we have are so many years of existence.

In the end it might be a case where one cannot choose whether one is an episodic or a diachronic narrativist any more than they can choose their basic personality or temperament. And this may in turn influence the badness of their own death, how much harm it does them.

4.3 Conclusions

Developing life extending technology to prolong the human lifespan is a worthwhile pursuit for the reason that death can be a bad thing for the one who dies, making postponing death, ceteris paribus, a good thing. My deprivation event view of the badness of death demonstrates this, while at the same time avoiding problems faced by Nagel's deprivation state view of the badness of being dead. Obviously not everybody will desire more life in cases where death is not a bad thing for the one who dies, but the deprivation event view of the badness of death allows for these cases. But by virtue of

this account we are rationally obligated to prolong lifespans in cases of premature deaths.

CHAPTER 5

LIFE EXTENSION AND FUTURE GENERATIONS

> Perhaps what we should do is be guided by concern for what will lead to the greatest total amount of happiness, or welfare, over time. On this basis, we should recommend against any further development of the anti-aging drug.
>
> Peter Singer, 'Research into Aging' 1991: 144.

The combined conclusions of the previous two chapters are that not only is the desire for profound personal survival rational, but avoiding death, all else equal, is a good thing. I am going to assume the truth of these two conclusions moving forward. It should be uncontroversial now to say that life extension would benefit the individual who wishes to prolong his or her own lifespan. However, an individual life does not happen in a vacuum, and the fact that people start having radically longer lives is going to affect more than just the individual. For example, one of the most common questions asked about life extension is whether allowing people to extend their lifespans will negatively affect the wellbeing of existing people and of future generations. If this were so, would it then be impermissible to develop life extension? In this chapter I will start to give my answer to the normative question of the thesis, 'Should we develop life extension', by first answering the challenge that developing life extending technology would be wrong.

In an article entitled 'Research into aging: should it be guided by the interests of present individuals, future individuals, or the species?' Peter Singer (1991) argues that developing life extension would not be in the interests of future generations. He accepts as uncontroversial that life extension would benefit present individuals who wish to extend their lifespan, other things being equal. And he is agnostic about whether life extension would be good for the species as a whole, indeed he raises doubts about whether the question is even intelligible; after all, can a species have interests apart from the interests of individual members of the species? However, Singer recognises that the more perplexing question is whether developing life extension would be morally permissible, based on how this decision would affect future people, and he offers us a hypothetical scenario, which I illustrate below, in order to demonstrate why

this question is so complex. Singer's own conclusion, based on his hypothetical scenario, is that we should reject a life extending drug because developing it would result in a world with lower total and average happiness. I will argue, contra Singer, that developing a life extending drug would not be impermissible. As we will see, Singer's hypothetical scenario is noticeably sympathetic to arguing that life extension leads to lower total and average wellbeing. The case is loaded in his favour. Therefore, if I can demonstrate the permissibility of developing life extension within Singer's scenario, then this only strengthens my general case.

First I will describe Singer's life extending drug scenario and the puzzle it presents us. Then I will evaluate Singer's recommendation in the case. He insists that we be guided by either a total or average consequentialist principle in our decision whether or not to develop life extension, both of which, in the end, supply us with the conclusion that we are better off rejecting it. I will suggest an alternative formulation of the average principle, under which we are presented with a different conclusion about developing the life extending drug, and I will offer an independent argument that we would be better adopting this consequentialist principle over others.

5.1 Singer's life extending drug

A new drug has been developed that dramatically slows human ageing, effectively doubling the lifespan to about 150. Taking effect in middle age, people who opt to take the drug enjoy an additional 70 to 80 years beyond a normal maximum lifespan, the traditional three score and ten. Importantly, however, the quality of life of those extra years is not quite as good as the average quality of life enjoyed over a normal lifespan.[32]

If we choose to develop the drug, says Singer, we must also cut the number of births, because doubling the lifespan will otherwise lead to more people existing at any one time than the world can support. One way of doing this is that we persuade people to have fewer children by procreating later in life, perhaps when they reach 50 instead of when they are about 25. Women on the drug, who are 50 years old, will still have the reproductive health of the average 40 year old, and will for many decades, making them

[32] Singer's constraint on the nature of the drug is not unrealistic. Since we do not know what effects radical life extension, whatever form it takes, will have, it is perhaps a safe assumption that a life extension drug will increase the user's maximum lifespan at the cost of at least some quality of life.

biologically capable of having children later in life. Thus after about a century or two of people using the life extending drug and having babies less frequently it will be true that fewer people will have lived than if we had rejected the drug (Singer 1991).

We have to assume here, on Singer's behalf, a finite career to the human species. To say that humanity and the earth will not last forever, given the limitations of the physical world, might seem a truism. But an end point to the existence of humanity is implied by Singer's claim that fewer people will have ever existed if we developed a life extending drug and lowered the rate of reproduction, an eventuality, as we shall see, that contributes to the total amount of happiness in the outcome. So for the sake of simplicity, assume that, whether we develop the drug or not, humanity will last strictly for another 10,000 years, perhaps because 10,000 years from now a rogue planet crashes into Earth annihilating all life instantly.[33] If we grant this, it will be true that far fewer people will have ever existed if we develop the life extension drug and procreate less, than if we had rejected the drug.

Given all of this, should we develop life extension?[34]

Fig. 5.1 The life extending drug

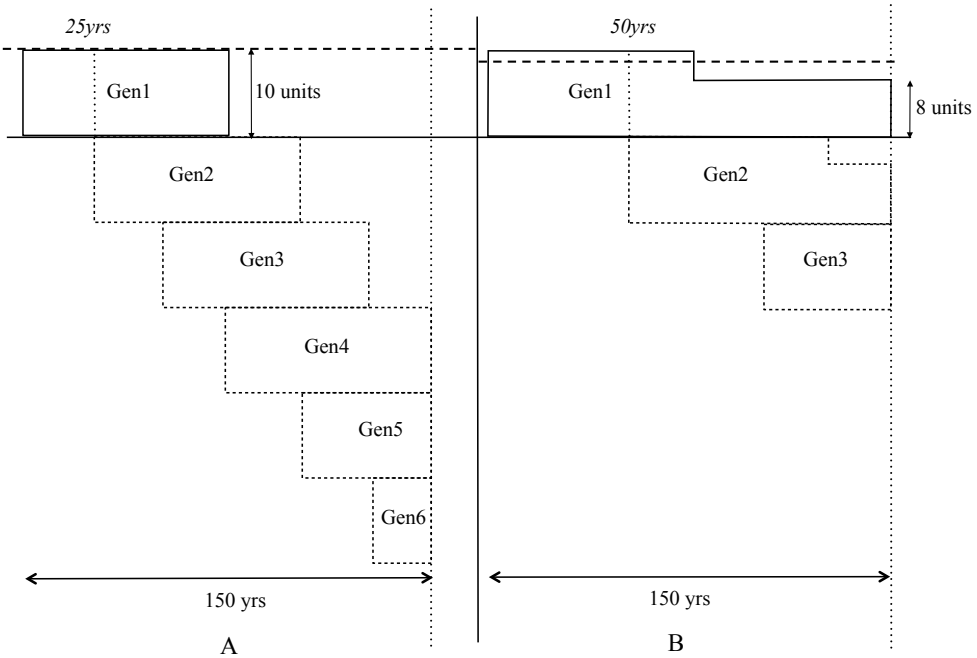

[33] *Melancholia* 2011, Lars von Trier.

[34] Singer's scenario is, by his own admission, a hypothetical one. We can assume that he is drawing conditional conclusions about whether developing a life extension drug would be the right thing to do based on reasonable assumptions about what effects the drug might have. Thus, his argument doesn't depend on the plausibility of his scenario, which, in any case, is something we have no empirical data for anyway. Attacking the plausibility of Singer's scenario has been a large part of the approach of two prominent responses to Singer's article: Blackford 2009; Walker 2006.

Figure 5.1 is a detailed representation of Singer's hypothetical scenario. There are two possible worlds we can choose to bring about, world A, in which we reject the drug and replacement (procreation) is normal, or world B, in which we develop the drug, and people have lives twice as long, but replacement is less frequent. Whichever world we choose to bring about will be the actual world.

So what exactly does figure 5.1 show? One of the first things to say about it is that it vindicates Singer's claim that fewer people will have ever existed if we developed the life extending drug. Each block represents a person's life. Six people, or generations, exist in A, they have children every 25 years, while three generations exist in B, and they procreate only every 50 years. The width of the blocks represents the length of a person's life, with the people in A living for 75 years and those in B for 150 years. The height of the blocks represents the level of wellbeing at any given moment within a life. This, of course, is an idealized representation of wellbeing across a life. In actuality it would go up and down allowing for the vicissitudes of individual happiness at different times during their lives. For the people in A, not on the drug, wellbeing is 10 units across their entire lives. For the people in B, who have taken the drug, it is 10 units for 75 years, then 8 units once the drug has kicked in for the remaining 75 years of their life. As stipulated by Singer's scenario, the extra years experienced by people in B are at a level of welfare not quite as good as that enjoyed over a normal lifespan. This is indicated by the drop in wellbeing during the extra 75 years lived by people in B.[35] On the other hand there is no dip in wellbeing over the span of the lives of those in A, who rejected the drug.[36]

As a result, the average happiness at any given moment rounds off slightly lower in world B than in world A, as indicated by the broken line in both outcomes. The total area in each block represents the total welfare per person; we can see there is more wellbeing *per life* in B. And finally, the total area of the blocks in each world represents the total welfare over time in each of the two worlds, and thus there is also more total

[35] There is a small inconsistency here in Singer's description of the case. The life extension drug kicks in after the person in world B reaches middle age (approx. 40 years old) and she continues to live at the biological age of a 40 year old for the rest of her life, but the extra years she lives over and above the regular three score and ten are, Singer stipulates, lived at a lower quality of life. So the mysterious dip in quality of life occurs, not when the drug starts to do its work at around 40, but only after the person has reached a 'fair innings', after which wellbeing is slightly lower for the remaining 70 or 80 years.

[36] The level of welfare in outcome B will still fall just slightly lower to that of A, now matter how many generations overlap at any one time, so long as at least one generation experiences the dip in wellbeing on the life extension drug. The average happiness at any given moment in A, on the other hand, remains unchanged, even given the introduction of new generations at any given time.

welfare in A than in B, because there are more people existing in world A in which we reject the drug. In B we've brought about a world where we've removed extra lives and replaced this with more life per person at a level not as good.

5.2 Total and average utilitarianism

Now that we have Singer's puzzle before us, how do we decide which world would be better: should we reject the life extension drug or develop it? There are two obvious normative principles that we might rely on to guide our decision, the 'total principle', which says it is permissible to bring about an outcome if and only if it contains no less total happiness than any alternative outcome; and 'the average principle', that it is permissible to bring about an outcome if and only if the average happiness in that outcome is no lower than in any alternative outcome.

On the one hand we have two claims that developing a life extending drug would be impermissible, because it will result in less total happiness and lower average happiness at any given moment. And on the other hand we have the intuition that extending people's lives is a good thing, all else equal.

The total view is fairly straightforward and unambiguous. It tells us that the best thing to do is to maximise the total happiness over time for any possible outcome. Thus, in fig. 5.1, for example, we simply sum the happiness for all lives over all time. And if this measure is higher for world A than the measure for world B, then it is obligatory to bring about world A.

The average view, however, can be interpreted in different ways, with different renderings of the principle yielding different permissible choices for the same case. Singer does not, nor do many other philosophers, distinguish between different formulations of the average principle.[37] Parfit has given the average principle as "If other things are equal, the best outcome is the one in which people's lives go, on average, best" (1984: 386; 420);[38] or negatively "it is worse if there is a lower average

[37] Hurka 1982 is one of the only exceptions in the literature. He formulates six different versions of average utilitarianism. However ingenious Hurka's different readings of the average view, many of them are found to be, in essence, hybrid principles, making use of both total and average calculations. Therefore I would argue that he has not sufficiently captured the spirit of what an average principle is in the majority of his different interpretations.

[38] This is Parfit's impersonal formulation of the average principle. He also has a person-affecting version of the average principle that includes the person-affecting restriction, i.e. that an outcome is best

quality of life, per life lived." Rawls formulates the principle of average utility as that institutions should be arranged so as to maximise possible average happiness (1971: 161-162). And Smart has expressed the average view as the view that "we should try to maximise the average happiness of human beings (or the average goodness of their states of mind)" (1973: 27).

These are all rather scant descriptions of a major normative principle, and Singer's formulation is no different. Singer states the average view as the view that it is best to bring about an outcome where "the average level of happiness at any given moment" is higher (1991: 140-141). The simplest interpretation of what Singer means by this is: The average level of happiness in a world is the amount of happiness one would expect to find in the world at a randomly selected moment in time. To calculate this, we divide the total amount of happiness in a world by the total number of times, for as long as people exist, in that world. I will call this the 'average happiness per moment principle'. Furthermore, I would suggest that this is the orthodox understanding of average utilitarianism.

I, on the other hand, will defend a different interpretation of the average utilitarianism: the average happiness in a world is the amount of total happiness one would expect a randomly selected *person* in that world to have. To calculate this, we divide the total amount of happiness in a world by the total number of people in that world. I will call this the 'average happiness per person principle'.

Under the total principle and the average happiness per moment principle, we choose best by bringing about world A and rejecting the life extension drug; but under the average happiness per person principle, we choose best by bringing about world B by developing the drug. To see why this is, consider the following calculations. The application of the three consequentialist principles above to the scenario shown in fig. 5.1 are as follows.

(1) *The total principle*: it is permissible to bring about an outcome if and only if it has the greatest total amount of happiness over time than any other outcome. So if we sum the happiness for all lives over all time in fig. 5.1, we get:

---

because it is best *for* the people who exist, because the lives of these people are on average better. So it is the people, not the outcome, that are the focus of the goodness or badness of our choice.

A: Gen 1 (10 x 75) = 750 +
Gen 2 (10 x 75) = 750 +
Gen 3 (10 x 75) = 750 +
Gen 4 (10 x 75) = 750 +
Gen 5 (10 x 50) = 500 +
Gen 6 (10 x 25) = 250

A = 3750

B: Gen1 (10 x 75) + (8 x 75) = 1350 +
Gen2 (10 x 75) + (8 x 25) = 950 +
Gen3 (10 x 50) = 500

B = 2800

As we can see there is quite a bit more total happiness over time in outcome A than in outcome B.

(2) *The average happiness per moment principle*: it is permissible to bring about an outcome if and only if the average happiness per moment in that outcome is no lower than in any alternative outcome. For the sake of practicality, if we determine the level of happiness at, say, any given year that people exist in each outcome in fig. 1, we get:

A: Gen 1 (10 x 75) = 750 +
Gen 2 (10 x 75) = 750 +
Gen 3 (10 x 75) = 750 +
Gen 4 (10 x 75) = 750 +
Gen 5 (10 x 50) = 500 +
Gen 6 (10 x 25) = 250

3750 ÷ 150

A = 25

B: Gen1 (10 x 75) + (8 x 75) = 1350 +
Gen2 (10 x 75) + (8 x 25) = 950 +
Gen3 (10 x 50) = 500

2800 ÷ 150

B = 18.6

The average level of happiness in outcome A rounds out only slightly higher than in outcome B. The levels of happiness 10 and 9.3 correspond to the broken lines in A and B respectively, shown in fig. 5.1. This result once again proves Singer correct in his conclusion about developing the life extension drug.

What result do we get in fig. 5.1 if we were guided by the alternative average principle I proposed? If we apply the average welfare per person principle to case we get a different answer than the other two principles.

(3) *The average happiness per person principle*: It is permissible to bring about an outcome if and only if the average happiness per person in that outcome is no lower than in any alternative outcome. We first determine the total amount of happiness in a world and then divide by the total number of people in the world to determine the average amount of happiness per person:

A: Gen 1 (10 x 75) = 750	B: Gen1 (10 x 75) + (8 x 75) = 1350
Gen 2 (10 x 75) = 750	Gen2 (10 x 75) + (8 x 25) = 950
Gen 3 (10 x 75) = 750	Gen3 (10 x 50) = 500
Gen 4 (10 x 75) = 750	
Gen 5 (10 x 50) = 500	
Gen 6 (10 x 25) = 250	
3750 ÷ 6	2800 ÷ 3
A = 625	**B = 933.3**

The result is a much higher average happiness per person in outcome B than in outcome A. Under this third average principle, it is permissible to bring about world B by developing the life extension drug, because in doing so we are bringing about the outcome with the greatest average happiness per person.

How do we decide which of the three principles we should be guided by in the life extension case (and, perhaps, in other population cases)? In sections 5.3 and 5.4 below, I will first show that, neither the total principle or the average happiness per moment principle are able to provide sufficient conditions to show us which world we are obligated to create while at the same time avoiding their negative implications. Then in section 5.5, I will present an independent argument for preferring the average happiness per person principle to the average happiness per moment principle.

5.3 Problems with the total principle

5.3.1 The open ended future of humanity

Earlier I observed that the claim that there will always be greater total happiness over time in the world in which we reject life extension can only be true in virtue of there being a foreseeable limit to humanity's survival. It would be interesting, then, to see what would happen if we alter Singer's scenario slightly in this regard.

Let's assume that the rogue planet doesn't crash into Earth 10,000 years from now and all life is not instantly annihilated. Given an open ended future, will it always be true that fewer people will ever exist on the life extension drug even if we have children less frequently? It might be argued that developing the drug and ushering in an era of prolonged lives over a long period of time would have an affect on the longevity of the species as well. For example, say in 150 years we finally reach the limit of possible reproduction cycles available to human beings and can no longer make babies.[39] If this were to happen, life extension may save humanity from a more immediate extinction. This is because artificial life extending interventions, like the life extension drug, will be left unaffected by these reproduction related population-limiting factors. In this case, people will have existed for a longer period of time in the world in which we developed the life extension drug than the one in which we rejected it, resulting in the greater total sum of happiness over time.

[39] *Children of Men* 2006, Alfonso Cuaron.

Fig. 5.2 Humanity's open ended future

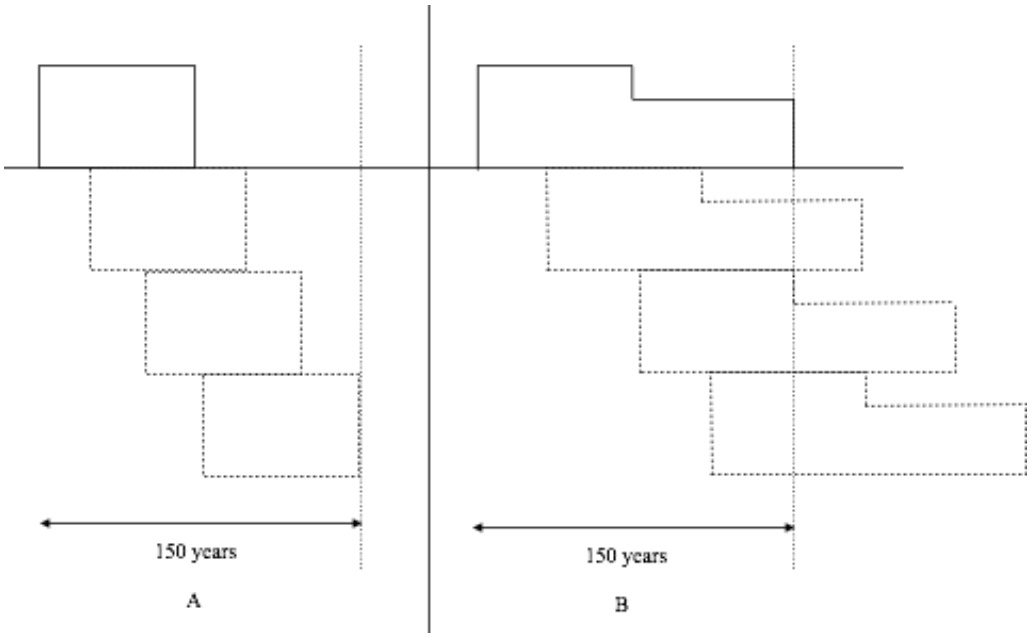

As in fig. 5.1 in fig. 5.2 above each block represents a person's life, the length the person's lifespan, and the height their level of happiness. In both outcomes A and B reproduction is terminated after four generations, and thus humanity is cut off after 150 years in world A, where the life extension drug had been rejected, but continues for a while longer in B on account of longer lifespans. We can see that the total area of all the blocks is now greater in B than it is in A, so that world B contains greater total happiness over time than A.

Of course, neither Singer nor we can know the contingencies of humanity's eventual extinction. Nonetheless it is worth highlighting this potential variation on Singer's original scenario, in which rejecting the life extension drug in favour of procreation increases the likelihood of shorter existence for humanity. However, for the remainder of the chapter, I will presume that more people will have ever existed in the world in which the life extension drug is rejected.

5.3.2 The repugnant conclusion

As consequentialist theories, the total and average views say that all we ought to care about is the goodness or badness of the outcome resulting from our actions or choices.

In their utilitarian form, what is good or valuable that outcomes must maximise is happiness. Thus to be guided by the total principle we are committed to bringing about that outcome with at least as much happiness as any other available outcome.

Fig. 5.1 shows that rejecting the life extending drug will lead to more people living happy lives, and therefore a world with greater total happiness over time. This is indicated by the total area of the blocks that make up world A compared to the total area of blocks that make up world B. Extending people's lifespans, it seems, is not an optimal way of maximising total happiness over time. Bringing new people into existence, so long as they have a worthwhile life and the extra happiness brought about by every additional person is greater than the unhappiness suffered by those already existing persons as a result of bringing a new person into existence, is the best way to maximise total happiness in the world. Consequently, one of the implications of the total principle is that we should strive to bring as many new people into existence as possible, so long as their life contributes to the positive welfare. This principle leads to some unacceptable implications, chief among which is the notorious repugnant conclusion (see: Parfit 1984: 381-390).

The total utilitarian must eventually accept that, for any world populated by individuals who enjoy a very high quality of life there is always a better world, one that we should choose to bring about, with a much larger population of individuals, all with lives that are barely worth living, because the total happiness in this world is greater.

Fig. 5.3 The repugnant conclusion

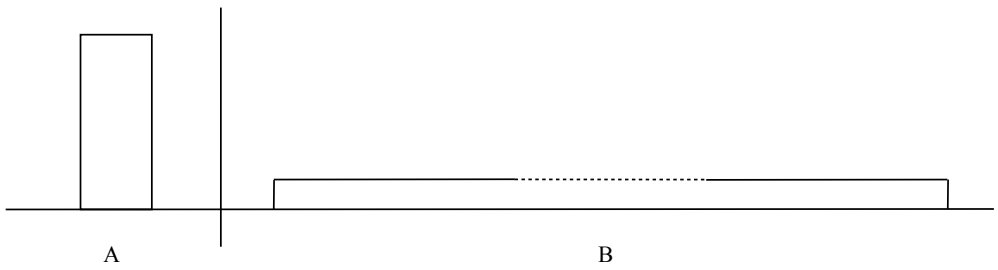

Fig. 5.3 is a depiction of the repugnant conclusion. The width of the blocks is the size of the population while the height is the level of happiness enjoyed by individuals. The population that makes up world B is vastly larger than the population in A but people in

B experience a greatly reduced quality of life. Despite this, world B clearly contains more total happiness than world A, as indicated by the area of the blocks. Therefore, we should prefer world B if we were guided by total utilitarianism. Most people find this implication of the total principle hugely unattractive.

5.3.3 Voluntary extinction

To be guided by the total principle is to be committed not only to the happiness of the people who exist or will exist independent of our choice, but, as Singer says, to the interests of 'merely possible' people as well (1991: 144). Or put another way, total utilitarians must not only be impersonal in their moral decision-making but also adopt temporal and modal impartiality. To illustrate this duty to the interests of merely possible people, Singer uses the familiar case of the choice the present generation must make between depleting the world's resources or conserving them for future generations.[40]

Briefly, as a global community we must make a choice whether to deplete or conserve resources. If we choose depletion, the quality of life will be higher in the immediate future for existing people, but much lower for future generations after the resources give out. On the other hand, if we choose conservation, then the quality of life in the immediate future will be slightly lower, as we struggle to find alternative forms of resources to replace the resources that could have been consumed, but unconsumed resources will be available to future generations, ensuring their sustained quality of life.

Singer suggests that one possible solution to the conservation problem open to the present generation is voluntary extinction. We could simply decide to become the last generation on Earth. There would be no more population growth and we could conceivably consume and pollute to our hearts' desire without having to worry about any moral questions relating to the happiness of future generations. Singer finds this choice impermissible.[41] Those guided by the total principle must reject outright the

[40] See: Parfit 1984: 361-364.

[41] Singer has asked the question of whether it would be wrong for us choose to become the last generation on Earth in a recent piece in The New York Times Opinionator, Singer 2010. David Benatar, in his book *Better Never To Have Been*, 2006, argues that it would be better if humans gradually became extinct, preferably through a voluntary end to procreation, because of the suffering involved in human existence. There also exists a real world movement, called the "Voluntary Human Extinction Movement" (VHEMT) <http://www.vhemt.org/>, whose aim is to promote the gradual, voluntary extinction of the

permissibility of voluntary extinction mainly because we would not be maximising the possible happiness over time. Moreover, many people might intuitively agree that voluntary extinction is impermissible. However, I would argue that the permissible choice between replacement and conservation, on the one hand, and voluntary extinction and consumption, on the other, might be voluntary extinction. This is because, if we are guided by the total principle, we once again run up against the repugnant conclusion.

Fig. 5.4 Voluntary extinction

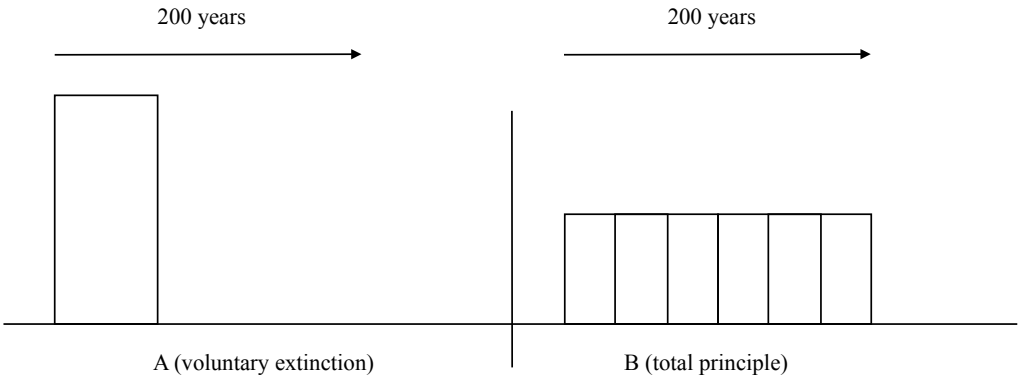

Consider the figure above. The rogue planet hits Earth 200 years from now. Which world would it be better to bring about? Outcome A, in which the present generation decide to have no more children and become the last generation on Earth? In this world we need not concern ourselves with conserving resources for posterity but can live out the rest of our lives (which is also the balance remaining of the career of humanity, though not yet of the earth) with a high level of wellbeing.[42] Or would it be better to bring about outcome B, taking as our moral guide the total principle, where we are obliged to have as many children as possible while at the same time conserving resources for future generations because this will maximise the total possible happiness over time?

Outcome B does in fact contain a greater total amount of happiness than A. However, outcome A, voluntary extinction, results in much higher average level of welfare, both per moment (at least during the period of time that humans beings exist in

human species through non-reproduction, but unlike Benatar, their motive for this is apparently in order to return the Earth to its former glory.

[42] The high level of welfare would obtain at least until the point where there are very few people left in the world, whereupon quality of life would dramatically drop.

the world, though maybe not if the average were calculated over the full 200 years, when the world in inhabited and when it is deserted) and per person, for the last generation on earth, and no future generations, thus no people who will exist in the future.

If we compare the scenario in fig. 5.4 above with that fig. 5.3, which depicts the repugnant conclusion, we see that they are almost identical in their choice between potential outcomes, and we should also I would argue be the same in our intuitive response to them. Just like in the repugnant conclusion scenario in which we should prefer outcome A with the smaller, happier population, to the vast population of miserable people, we would likewise prefer voluntary extinction, if we believe it would be better if existing individuals have happier lives and are neutral about the fact that there would be no new people brought into existence. Singer, in arguing the impermissibility of voluntary extinction and the obligation to replace and conserve, is, unknowingly, arguing for the repugnant conclusion.

5.4 Problems with the average happiness per moment principle

Even if we reject the total principle as guide in our choice of whether to develop the life extending drug, it remains the case that developing the drug would result in a world with lower average happiness per moment, which may appear to constitute another reason to reject the drug.

Though the average happiness per moment principle avoids the repugnant conclusion (fig. 5.3) as it clearly implies that it would be better to bring about world A where people enjoy lives of higher level of happiness, it has other counterintuitive implications.

Fig. 5.5 God's choice

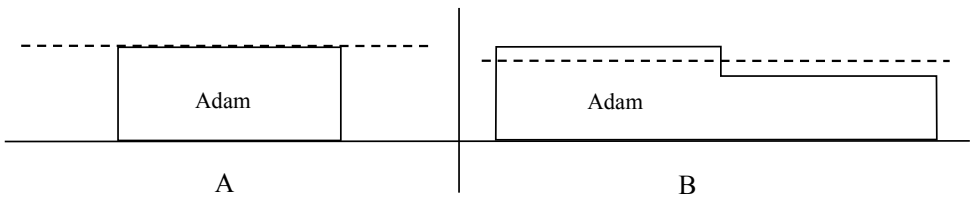

Consider God's choice whether to create world A or world B. In both worlds one person exists, Adam. In world A, Adam has a normal lifespan, and in world B, an extended lifespan. In world B, although Adam enjoys a longer lifespan, the outcome contains lower level of happiness per moment than world A, as shown by the broken line.[43] God's decision to make B instead of A will benefit Adam in one sense, namely, he will live a longer life than he normally would have, but he would be creating a world with a lower level of happiness per moment. Thus, if God were a regular average utilitarian, he would be making the superior decision in creating world A, by conferring on his Adam the normal three score and ten and denying him knowledge of how to make life extension.

God is acting rightly, under the principle that it is permissible to bring about an outcome if it contains the highest level of happiness per moment by bringing about A and not B, and he would be acting sub-optimally by bringing about B and not A. But what if we alter the above case slightly, and assume that God has brought about a different outcome, call it world C. There are two people in C, Adam and his twin brother, also named Adam. God grants one brother, 'Adam B', life extension, but not 'Adam A'. God then realises that he's made a mistake because the non-uniform happiness of Adam B's life is bringing down the average happiness at any given moment in the world (something which God, unfortunately, forgot to calculate beforehand). Thus, God, under the standard account of the average utilitarian view, would be quite just in killing Adam B, other things being equal, in order to maximise the highest level of happiness at any given moment in the world.

It hardly needs to be said that God's 'corrective' manoeuvre here would conflict with most people's intuitions. Surely it would not be better to kill Adam B because his life, which slightly sub-average in happiness in certain parts, is bringing down the average level of happiness at any given moment in the world. The tension, I think, lies in being committed to thinking about happiness in terms of an abstract level of happiness per moment, instead of thinking about the happiness within lives; the latter is still consequentialist thinking, but it is not a 'temporal' calculation, rather an 'accumulative' one, focused on persons. What the average happiness per moment principle does, it seems, is ignore a potentially morally relevant factor in our decision

[43] I am incorporating Singer's constraint about the lower quality of life of the extra years enjoyed by virtue of the life extending drug.

making in population cases, that is, the happiness accumulated in an individual's life taken as a whole. Figure 5.5 shows how Adam in world B has a better life than he otherwise would have, in the sense that it contains a greater amount of happiness, due to its longer length, than Adam's life in world A.

5.5 Arguments for the average happiness per person principle

Before I present an independent argument for preferring the average happiness per person principle over the total and average happiness per moment principles, I want to first see whether it manages to avoid the same negative implications that I charged the average happiness per moment principle with in the last section, namely, that the principle implies we create sub-average lives. Does the average happiness per person principle imply that we should not create sub-average lives? Recall the principle says that, it is permissible to bring about an outcome if and only if no alternative outcome has greater average happiness per person. If God were guided by this principle in fig. 5.5, he would do best to bring about world B which contains greater average happiness per person.

But what about in my alternative description of the case, does the average happiness per person principle imply God would do best to kill an individual whose life contains less than the most happiness per person? No it does not. In this version of the scenario there is one world containing two people, Adam A, with a normal lifespan, and Adam B, with a longer lifespan. Adam A is not bringing down some global level of happiness by his shorter existence, but he is contributing less than optimally to the average happiness per person in the world. But the average happiness per person principle does not imply that God would do best to euthanize Adam A, quite the opposite, it says that God would do best to give Adam A *more* life, to prolong his lifespan.

5.5.1 An argument for the average happiness per person principle

The result we get from being guided solely by either the total or average welfare per moment principles is that we would be better off rejecting life extension, which is a

result, I believe, that goes against commonsense intuition in the life extending drug scenario. Given the basic choice between extending the life of an existing person who wants to continue living, or bringing a new person into existence, I think that most readers would say that we should extend the life of the existing person if this makes their life better, and that we do not have an obligation to bring new people into existence to increase the average level of happiness.[44] Might it be, then, that the average happiness per person principle better captures our intuitions about the value of extending or improving the lives of existing persons? This principle takes into account the fact that, in fig. 5.1 for example, the lives of the people in B, which are happy lives throughout, are twice as long as the lives of those in A, therefore individual persons are experiencing much more happiness, and in this sense have better lives.

That being said, if I want to avoid begging the question by merely claiming the intuitive superiority of the average happiness per person principle over the average happiness per moment principle, I need to give an independent argument for preferring the former. Such an argument could involve the following. Singer is committed to caring about happiness per moment in the world not strictly the happiness of persons, but think about a scenario in which we have two possible outcomes we can bring about, world A, in which three people exist at the same time for 100 years and then no one exists for 200 years, and world B, in which three people also exist but only one person exists at a time, spread out over 300 years. Fig. 5.6 depicts this scenario.

[44] This basic attitude has been referred to as the 'asymmetry' in the literature, though note that the version of the asymmetry I have formulated above is not quite the standard formulation, which is that we seem to think there is a moral difference between 'making people happy' and 'making happy people'. See: McMahan 1981; Narveson 1967. See also Broome 2005: 400 for discussion of this 'basic intuition'.

Fig. 5.6 Average happiness per person

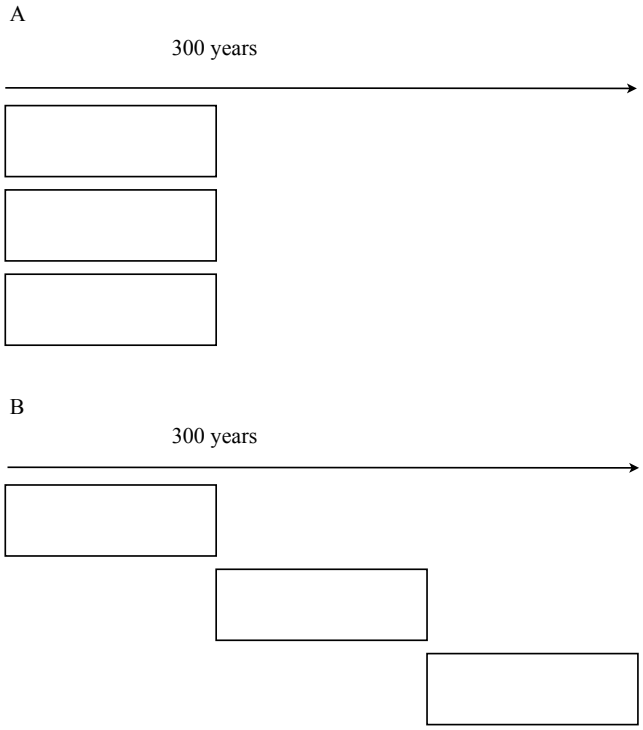

The total happiness in each outcome is the same and the number of people existing in each outcome is the same. Also, the average happiness per person is the same in both outcomes, since we simply determine the total amount of happiness in each world and divide by the total number of people in each world. However, there is higher average happiness per moment in A than in B. This is because, while the total happiness is the same in both worlds, the time that people exist is shorter in world A, the result being that there is higher average happiness per moment in A.

Consequently, those guided by the average happiness per moment view must conclude that the only permissible outcome to bring about is outcome A. But this result doesn't seem plausible. Surely it is no worse to bring about world B, in fact, it might in one sense be the better outcome, better in the sense that people exist for a longer period of time, i.e. 300 years instead of only 100 years. Those guided by the average happiness per person view can say that it is permissible to choose either world. The average happiness per person principle says that average happiness in a world is determined by the total happiness in the world divided by the number of people existing in the world, and since total happiness and number of people are the same in both outcomes, it is permissible to bring about either one. This result, I argue, is more plausible.

The average happiness per person principle also says it is permissible to choose the outcome in which we develop life extension in figure 5.1, because if humanity opted for the life extension drug (world B) individuals will live longer lives than those in world A, resulting in greater average happiness per person in world B. Notice, moreover, that the average happiness per person principle says it is permissible to choose voluntary extinction in figure 5.4 because individuals in the voluntary extinction outcome enjoy both a higher level of average welfare at any given moment and average welfare per person over the course of their lives. In each of these cases, the principle favours outcomes that, I would argue, are both the most intuitively plausible, and in which individual lives are better off.

5.6 Conclusions

I have argued that Singer does not choose the best moral principle to decide the case of whether to develop life extension and as a result favours the least permissible outcome, namely, that we should reject a life extending drug. I have suggested that a better consequentialist principle to use is what I have called the average happiness per person principle. If we are guided by this moral principle we get the conclusion that it is permissible to develop the life extension drug and prolong people's lifespans. The average happiness per person principle also avoids certain counterintuitive implications that the total principle and average happiness per moment principle are susceptible to. In Chapter 7 I explore whether we get a similar conclusion if we are guided by nonconsequentialist principles.

In Chapter 3, I concluded that profound personal survival is not necessarily a bad thing, and in Chapter 4 I concluded that death is a bad thing for the one who dies, other things being equal. In the present chapter I have concluded that, if we were guided by a plausible consequentialist moral principle, developing life extension would be permissible. The combination of these conclusions will allow me to assume in the following chapter, where I move to examine what justice says about how we ought to distribute the life extending technology we develop, that life extension is a good thing to have but that we must now decide on the fairest way to distribute it.

CHAPTER 6

LIFE EXTENSION AND JUSTICE

> ...if immortality or increased life expectancy is a good thing, it is doubtful ethics to deny goods to some people because we cannot provide them for all.
>
> John Harris, 'Intimations of immortality' 2002: 8.

In Chapter 5 I offered an answer to the question of whether we should develop life extension that employed exclusively consequentialist principles. In giving the answer that developing life extension is permissible I assumed that what we ought to do about life extension depended upon which of our options would have the best consequences. However, some may question whether consequentialist moral principles are our best guide in thinking about the value of life extension, or whether being guided by these principles preclude us from morally better answers.[45] By applying consequentialist moral principles we identify an action—for example, developing a life extension pill—that produces better consequences than any other alternative action. But perhaps some actions are wrong or obligatory despite the value of their consequences?

If we could predict that life extending technologies would be accessible only to the wealthy and those directly involved in their development would it be wrong to develop the technology in the first place? Much of the discussion about life extension and justice has been framed around the wrong question, the question I have just stated. For example, Singer does this by the fact that his argument takes seriously that we should prevent life extending technology being developed in the first place if doing so would bring about the outcome with the greatest wellbeing.[46] But this is not the most relevant question to ask. It is inevitable that we will have life extending interventions of some form or another some time in the future. Because of this, the question then becomes: if life extension were eventually developed, what does justice say about how we should distribute it? Therefore, I think the more important question to answer is

[45] Russell Blackford suggests this in his 2009 article.

[46] Ackerman 2009; Farrelly 2007; Harris 2002; Horrobin 2005; Kass 2004; Pijnenburg & Leget 2007, all identify 'the justice objection to life extension' as variations on the objection that we should not develop life extension at all if it will largely only be accessible to the most privileged.

what does justice have to say about how we should distribute life extending interventions given their successful development. This is not an objection to life extending technology as such, but an inquiry into what the fairest distribution of it would be. This is the question I will address in this chapter.

My analysis of this question will proceed with the help of the story begun in Chapter 5 about the development of a life extension pill. Continuing this scenario, I will speculate on the early distribution of the life extension drug and how this might play out and change over time. I will apply several major theories of justice to this story and try to find which principles of justice best compliment our intuitions about how we should to distribute the life extension pill. Thus, the methodology for analysing this question of justice will be through the framework of the major theories of justice. These are libertarianism, utilitarianism, egalitarianism, and Rawlsian principles of justice, with the aim of seeing which theory, if any, helps us reach a reflective equilibrium between basic intuitions about what a just distribution of life extension would look like justified by a coherent moral theory.

The conclusion I arrive at by the end of the analysis in this chapter is that a modified Rawlsian view has many benefits as a just distribution of life extension. I am not claiming that non-Rawlsian theories of justice have, as a result, been refuted, or that I have demonstrated the superiority of the Rawlsian view over other theories of justice, merely that, in the context of how we ought to distribute life extension, my prioritarian version of the Rawlsian difference principle is a strong candidate for the just distribution of life extension.

6.1 The early distribution of life extension

Let’s continue using Singer’s scenario about the development of a hypothetical life extension pill from the previous chapter. A scientist announces that she has successfully developed a pill that doubles the average life expectancy from around 75 years to around 150 years. Following this announcement it is revealed that the work done by the scientist and her team of longevity researchers had been funded by generous private grants from wealthy individuals and organizations coveting any intervention for prolonging the human lifespan. As a result, the life extension pill begins its life as a

privately funded, developed, and owned resource. This new longevity technology unsurprisingly generates enormous interest, in large part from wealthy individuals in developed countries the world over, who desire to radically extend their healthy lifespan. Since the technology is not yet in the domain of government regulation and allocation, anybody with sufficient means (the pill is very expensive and must be taken regularly) can freely purchase as much life extension as they wish, constrained only by the economic resources at their disposal.

At this point in the career of the life extension pill not everyone can access it due to the drug being out of his or her reach financially. Naturally, objections are raised about the unfairness of the inequalities resulting from this free market arrangement of the life extension pill; after all a resource such as a pill that prolongs a person's healthy lifespan is of benefit to everyone, other things being equal. Those with anti-libertarian feelings argue that government intervention is necessary in order to redistribute the life extension pill justly. More people must have access to the technology than merely the most advantaged members of society, who, by virtue of having access to life extension, are widening their advantage over those less privileged.

I want to first examine the question of whether a free market arrangement in respect to the life extension pill would be wrong. That is, would it be worse to have life extension available only to a privileged minority than to not having life extension at all. I will look at potential extrinsic or relative harms resulting from inequalities in the distribution of life extension and attempt to show that appeals to these harms are insufficient in demonstrating the impermissibility of a libertarian distribution of life extension. Following this I will address the weaker claim that inequalities resulting from libertarian distribution of life extension would be unfair, or less than an ideal distribution. I will agree with this claim and argue that there may be a better or fairer redistribution of life extension derived from Rawlsian principles of justice.

6.2 Libertarian distribution of life extension

We have an intuition about justice that says that, so long as the things one owns have been justly acquired, then one is entitled to what one owns. The other half of the intuition is that it would be wrong, other things being equal, to restrict one's freedom in

the free exchange of resources that have been justly acquired or to suffer redistribution of what one owns. This principle of liberty is the primary value promoted by a libertarian theory of justice. The libertarian approach says that a just distribution of resources is simply any distribution that results from people's free exchanges, and that there is no moral reason for restricting people's liberty by redistributing resources, not even to correct involuntary inequalities, so long as resources are justly acquired (Nozick 1974).

The libertarian, therefore, would have no problem with the arrangement that has naturally formed around the newly developed life extension pill. Assuming individuals acquire the life extension pill justly, the demands of justice have been satisfied. What exactly does it mean to justly acquire the life extension pill? It means that the resources (e.g. money) with which the person paid for life extension were acquired justly, and the means with which they acquired those resources was just, and so on backwards through the individuals' commercial history (right back to the initial acquisition of resources in the state of nature). If these conditions were met, then the libertarian would find nothing wrong with a free market distribution of a life extension pill. Since inequalities of access to the new life extension pill are simply the result of a system of free exchange, with people making use of justly acquired and transferred assets, nothing unjust has taken place. The fact that not everyone can have immediate access to life extension is not within the scope of considerations of justice, according to the libertarian view.

However, would the libertarian distribution of the life extension pill necessarily be wrong if such an arrangement harmed those less well off who cannot access the pill? First we need to ask in what ways unequal access to life extension might harm those without access. One intuition is that those without access are harmed simply insofar as they are worse off than those with access to life extension. While such 'comparative' or 'extrinsic' harms are familiar, only this time the advantage would be in an area of human life that does not purely involve economic gain or equality of opportunity, but biological advantage as well (Horrobin 2005). In addition to their previous economic freedoms, those individuals or groups most privileged in society will be permitted greater lifespan, renewed health, a reduced window of suffering from age-related disease and debility, to name the major benefits. Is this purely relational harm resulting from the fact that some people have access to life extension while others don't make the inequality wrong?

Consider the example presented in the movie *Elysium* (Neill Blomkamp 2013), which portrays a divide between those who have access to life extending medical treatments and those less advantaged without access. In the movie, in the near future the most wealthy and advantaged members of society will live exclusively on an enclosed luxury habitation in orbit of Earth. These lodgings include immediate access to medical technology capable of curing all human injury and disease, for example cancer and damage caused by ageing, while the rest of society on Earth have access to only the most basic medical resources and are condemned to have short, unhealthy lives. Though fictional, the 'Elysium' example does strongly imply an injustice inherent in the situation of a divide between those with access to technology that would enhance healthy lifespan and those without. Beyond this inherent injustice, the downstream social and economic consequences of a 'longevity divide' between the healthy lifespan-enhanced and the non-enhanced would be that the most advantaged in terms of health and longevity, because of their early access to the life extension pill, would enjoy ever greater opportunity to accumulate wealth and power compared to those who had no access to the pill.[47]

Beyond the extrinsic, relational harms brought on by the fact that one is worse off compared to someone else, can we say that intrinsic harms result from the inequalities on access to life extension? Those without access to the life extension pill may suffer distress, a kind of psychological harm, as they watch the wealthy live life exempt from ageing and illness. There is also the suggestion that the prospect of life extension that could be enjoyed by people in the future makes the prospect of one's own death worse for those who will not have access to the technology, amounting to another kind of psychological harm for the have-nots (Davis 2004).

Consequently, the anti-libertarian might argue that these harms prima facie strengthen the case for equal access to life extension. But the libertarian might also accept this as a prima facie, but not an in principle, argument against libertarianism. Libertarians would agree that we only have the right to act freely, to a free exchange of resources, if this does not result in the harm of others. In other words, libertarianism is limited by the harm principle, the principle that preventing harm to others should restrict actions, but this should be the only restriction on one's liberty. But even if

[47] For a fictional example of this, in Bruce Sterling's novel *Holy Fire*, 1996, future societies are controlled by an established 'gerontocracy' who are wealthy and long-lived, following the invention of life extension.

libertarian free exchange mustn't cause harm we should not count the sorts of harm I described above. If it were always wrong to be in a better position than another person as a result of free exchange, then almost any exchange would be impermissible. For instance, the fact that those without access to life extension would envy those with access is not a particularly compelling objection to life extension. If we counted the psychological harm caused by envy as a restriction on people conducting free exchanges then the whole system of free exchange would break down. Furthermore, speculation about eventual harms resulting from a longevity divide is problematic in terms of existent empirical support for such claims. So in the end, the charge of psychological harm is not an objection to libertarianism in itself.

A sensible libertarian would reject the significance of either extrinsic or intrinsic harms resulting from free market set up. The worse that can be said about the situation arising from a libertarian beginning to the career of the life extension pill is that allowing those with the economic means freedom of access to life extension constitutes a Pareto-improvement on the state of affairs. That is to say, some members of society will receive great benefit from a particular resource without directly making other members worse off. Extra healthy life years are, other things being equal, a good, even when it is a good not shared by all who could or would want to benefit from it.

On the other hand, is there anything good that can be said about the fact that life extension is made available at all, even though it is only available to some? What about the potential trickle-down effects of the mere development of a life extension pill? Technology like the life extension pill should not only benefit those with direct access but may also have trickle down effects into a number of socially beneficial areas, particularly medicine and healthcare, in alleviating old age debilities and diseases, conditions that are an enormous drain on a society's economic and medical resources. Moreover, given that the life extension pill not only retards ageing but compresses the time in which the individual spends in morbidity at the end of their life, conceivably enabling them to work right up until this final stage of life, there might be much less need to save for things like retirement and old age for those on life extension. Consequently, more people will be able to do without superannuation because they can continue to work for decades. However, considering that the people using life extension

will mainly be the wealthy and rely less on saving money for retirement, this might not make such a huge difference.[48]

The scenario of the life extension pill might further play out like this: If life extending technology becomes progressively cheaper and more widely available as time goes on, then the interval between those who first have the opportunity to use the life extension pill and those to whom it becomes available down the track will be negligible. Because of the nature of the life extension pill (allowing one to age very slowly) those with initial access to life extension will over time differ little in the benefits they received from longevity from those who are granted access to the technology later on. In this respect life extending technology is different from a technology such as information and communication technology, which is continually updated and older versions made swiftly obsolete. Longevity is longevity whenever you start benefiting from it, you will make up the difference in the end. However, of course, some people who are very old at the beginning will miss out entirely on life extension because they died of old age before they were able to access it.

But this is not to say that everyone will eventually get access to life extension, just as information and communication technologies (ICT's) are getting cheaper and chapter, still many people are without access to the technology, for example, in the less developed countries of the world, until absent government intervention. But there is an analogy to be made between the distribution of new ICTs and new life extending technologies. We seem to have some precedence for accepting inequalities when it comes to the distribution of new technologies in the existence of 'digital divides' for example. A digital divide is a state of affairs in which some have access to ICT's, while others, less-favoured, do not. Although ICT's influence the social position and welfare of members within wealthy societies in which they are available, the fact that, globally speaking, individuals in other societies do not have access to these new technologies does not in itself affect their basic welfare. Consequently, some have argued (Rooksby & Weckert 2004) that the inequality that exists in digital divides are not inherently wrong but rather any objection on the ground of justice comes from how the divide may indirectly exacerbate certain social problems for those without access. We don't think

[48] This point will be brought up again below in my analysis of Rawlsian principles of justice and the distribution of life extension; specifically, one might argue that the money the government saves with fewer people relying on superannuation because of greater time actively working and earning an income be put towards the subsidisation of life extension, making it cheaper and possible for more people to be on the life extension pill down the track.

that computers should never have been invented, or life extension pill never developed, it simply raises challenges.

According to the libertarian, inequalities resulting from free exchanges are not unjust, but redistribution of resources from free exchanges is unjust. I agree with the first half of the libertarian principle, but not the second. Redistribution can be just, if it makes things *better*, namely fairer, than the libertarian setup. Thus the libertarian principle, or so I will argue, provides grounds for the permissibility of a pattern of distribution of life extension, but does not optimally satisfy the demands of justice in this regard. Kymlicka (2002) has argued that not only is the libertarian approach unfair but that, taken to its extreme, it is, in theory at least, a self-defeating theory of justice, because by rejecting any correction to or improvement of disadvantageous circumstances it is denying those less advantaged the ability of self-determination or liberty, the very principle it is based on. As a result, I would argue that a non-libertarian redistribution of life extension would better satisfy justice.

6.3 Government intervention in life extension

We now continue our story about the early years of the life extension pill. It has been five years since our scientist and her team made the life extension pill commercially available and during this time those with enough money have been buying and taking the pill. Of these individuals those who were beyond middle age at the time they began the drug regime have immediately experienced some rejuvenation to their health and vigour, reverting them physically and cognitively back to their former 40 year old self, while younger individuals who started taking the life extension pill have not yet experienced any real benefits from taking the drug but are comfortable in the knowledge that they will reap the benefits in health and longevity as they continue to live.

Though the libertarian would have no problem, other things being equal, if this free market arrangement were to continue, the anti-libertarian would surely be crying out for government intervention into, and redistribution of, the life extension drug, that something further needs to be done to satisfy the demands of justice. As I have argued, redistribution of life extension can be just if it makes things even better, namely fairer, than the libertarian arrangement. Therefore, I will assume that government intervention

is necessary because, even if no one is made worse off by unequal access to the life extension pill, it is still unfair that some have access to it while others, who desire to prolong their lifespan, do not.

The anti-libertarian voices have been heard, and after an initial five years trial period of the free market experiment, the government finally decides to step in and buy out and subsidize the life extension pill, making it part of the healthcare arrangement. The state, of course, does not have an unlimited amount of the still fairly expensive resource to distribute, at least not until more can be developed on a larger scale in government funded laboratories, thus a policy must be decided upon about how to fairly distribute the life extension pill. Should priority for life extension be given based on some objectively measurable variable, such as projected life expectancy, or maybe a factor like potential future contribution to society, i.e. projected social utility, so we are extending the lives of 'useful' people, or maybe a variable less objectively calculable, such as strength of desire or preference to continue living?

6.4 Utilitarian distribution of life extension

Aside from the intuition that we should distribute resources as equally as possible, there is a natural intuition that we should try to bring about the most benefit from our distribution of resources. The general utilitarian approach is to produce the greatest amount of goodness or utility, classically understood as pleasure or wellbeing, or on more sophisticated accounts, as satisfaction of preferences, in the outcome produced by our choices or actions. Utilitarianism, therefore, can be regarded as a theory of justice because it has implications about not only what is the right arrangement of social institutions, but what is the right distribution of resources, namely, one that produces the greatest amount of wellbeing.[49] In keeping with the formulation I gave to the total and average utilitarian principles in Chapter 5, the utilitarian principle of justice says that a

[49] I will once again use wellbeing as the good to be maximised when talking about utilitarianism as a theory of justice, though the 'preference' theory of wellbeing is often preferred in discussion of utilitarian as a political morality, for reasons that it is perhaps more precise to talk about maximising people's preferences for the good life and the things in it. The important thing to remember is that both wellbeing and preference satisfaction are wider notions than health or life expectancy, which are instrumentally good things in the context of this discussion of distributing life extension, and therefore encompass both notions of health and life expectancy.

social arrangement or distributive policy is permissible if and only if it results in no less than the greatest wellbeing of any alternative arrangement or policy.

How, then, could we distribute the life extension pill to maximise wellbeing in society? This might involve giving priority to those who could benefit the most from prolonging their healthy lifespan, such as people with low life expectancy, for example, other things being equal, it is better to increase person A's lifespan from 20 years to 80 years than to extend person B's lifespan from 80 years to 100 years. Or the utilitarian might maximise wellbeing by giving priority to those people who can best take advantage of extra healthy life years, for example, the wealthy philanthropist, who, in virtue of her vast economic resources, can with extra life years continue to alleviate the suffering of those less fortunate around the world. It is ultimately an empirical question, from which the utilitarian chooses which policy maximises the good. The utilitarian isn't worried about distributing the life extension pill as such; all they are concerned with is maximising wellbeing. The utilitarian, for example, might think that no policy of distribution maximises the good other than not having life extension at all, like we saw in the case of Singer's conclusion in the previous chapter.

So perhaps this is a good place at which to say something about what exactly the life extension pill is as a resource for the utilitarian. The life extension pill is a resource that promotes wellbeing by increasing an individual's healthy lifespan. This is because, even though people do not agree on what kind of life would be the most desirable, we can all agree that continued happy life is a good thing, a thing that, other things being equal, increases one's opportunity for good experiences, pleasure, or happiness, and also one's opportunity to contribute to society.[50] As a result, the utilitarian, in trying to maximise wellbeing through the distribution of life extension is, more often than not, maximising the number of healthy life years in society. However, to be clear, maximising the number of healthy life years in society does not necessarily lead to maximising wellbeing in society, for example, increasing the healthy life years of a sadist would arguably add to the disutility in the world; or, there might be someone, call them a 'utility monster' (Nozick 1974), capable of experiencing ten times the amount of pleasure per moment than a normal person, who if we had to choose between extending their life by 10 years and extending the lives of ten other people to the same degree, it would be better to extend the life of the utility monster.

[50] Of course, not everyone thinks this is true, for example, Kass and Williams don't think so. But I have already done all that I can to rebut them in Chapter 3.

The utilitarian approach to the distribution of the life extension pill will differ from another theory of the distribution of resources that is often compared to the utilitarian approach, the 'quality adjusted life years' economic approach to healthcare distribution, or QALY, which says that we should distribute healthcare resources in a way that produces the most healthy life years per dollar. I will discuss the similarities and further differences between the two theories later in the section.

Another thing worth remembering at this point is that the utilitarian is not interested in benefiting particular individuals, but only in there being more wellbeing in the world. As a result, depending on certain contingencies about the wellbeing derived from life extension by different types of individuals, such as the wealthy or 'utility monster', it seems that utilitarianism recommends some intuitively unjust distributive policies, for example, that what we ought to do is to distribute the life extension pill only to the wealthy minority if this results in maximising available wellbeing. I will argue that utilitarianism is, therefore, limited as a theory of justice and that we need some further principle to satisfy justice in how we distribute the life extension pill.

Moreover, if the utilitarian policy-maker maximises wellbeing, she will run into many of the same basic problems that I argued plagued the utilitarian consequentialist in the previous chapter. Maximising wellbeing through distribution of the life extension pill leads to the repugnant conclusion: if greater total wellbeing can be generated by distributing the life extension pill in such a way that results in an enormous population all with lives of minimal wellbeing as opposed to distributing life extension so that there is a smaller population whose individuals have longer lifespan and more wellbeing per individual, then, according to utilitarianism, the former is the policy we ought to enact.

Introducing the life extension pill might lead the utilitarian to the repugnant conclusion. Introduction of the life extension pill will likely increase population and thereby lower individual average wellbeing because of increased pressure on resources. Singer's move in Chapter 5 was to put a cap on population growth by limiting reproduction. If this is done—life extension is introduced and we have babies less frequently—then the life extension pill will decrease total wellbeing, as Singer shows. But Singer's move of reducing reproduction is actually a non-utilitarian move, because by artificially lowering the rate of reproduction we will fail to maximise total wellbeing. What Singer should have argued, as a total utilitarian, is that the life extension pill be introduced and we keep on having as many babies as we already do. This would

maximise total wellbeing but would lead ultimately to the repugnant conclusion. Contra Singer, utilitarianism is not against the life extension pill; it is against lower population.

There is another kind of problem for utilitarianism as a theory of justice. As I suggested earlier, individuals differ in the amount of wellbeing they get out of extra years of life, not to mention in their desire for continued life. In order to determine what kind of distribution would maximise the wellbeing derived from extra life years the utilitarian might take into account the fact that distributing the life extension pill to certain people will result in more utility than distributing it to others. Consider the hypothetical example of the 'utility monster'. Imagine there exists someone who would gain an enormous amount of utility from the same resources that would give most people only a small benefit. The utilitarian would have to say that an unequal distribution of life extension to these individuals would be the right redistribution because this resulted in the greatest overall happiness. There might be (though this is a fairly large empirical contingency) members of society who would be utility monsters of a sort when it comes to life extension, so that the benefit of allocating the life extension pill to these few people would outweigh that of distributing it to the majority. Perhaps the wealthy technophile or life extension fetishist, say, a Ray Kurzweil (1999; 2005; 2005b), would be representative of this kind of individual. If more healthy life years are especially beneficial to some individuals, then it seems that the utilitarian ought to prioritise distribution to such individuals because they can get the most benefit out of it.

If satisfying the desires of the wealthy, for example, to extend their lifespans produces the optimal level of wellbeing, then prioritising the extension of those people's lives might be the right policy under a utilitarian view. This is, of course, based on the empirical assumption that satisfying the preferences of the wealthy to extend their lifespans will generate extremely high levels of wellbeing enough to outweigh the wellbeing produced by redistributing the life extension pill more equitably. Maybe life extension ought to go to those individuals who really want a profound extension of their lifespan, as much as they can get. As I mentioned at the start of the section, this conclusion derived from the utilitarian principle of justice would likely run counter to many people's idea of a just distribution of a resource that increases a person's healthy lifespan such as life extension.

Surely the disutility resulting from those less well off being denied the opportunity to radically prolong their lifespan cannot be discounted so lightly. Even

factoring in to the utilitarian calculus the continued happiness accrued by the minority of privileged individuals who, since they are able to live much longer lives, will continue to contribute to the amount of utility in the world long after those who missed out on life extension are dead, it still seems that the disutility in the world resulting from the rest of the people who miss out on life extension would be unjust.

6.5 Quality adjusted life years approach

As I mentioned earlier, there is an existing economic strategy that resembles the utilitarian approach to justice. The quality adjusted life years approach, or QALY method, is a cost-benefit measurement for fairly distributing healthcare resources, a kind of economic rendering of the utilitarian theory. Under QALY, a particular allocation of medical resources is determined based on the number of quality adjusted life years (which is the number of years and the quality of those life years in a single measure) it will produce for every available healthcare dollar (Mckie *et al*. 1996). Thus it is a method for calculating the quality of life impact of a particular distribution of health care resources, where one expected healthy life year is worth one QALY. For example, say we need to choose between two lives, *x* and *y* below:

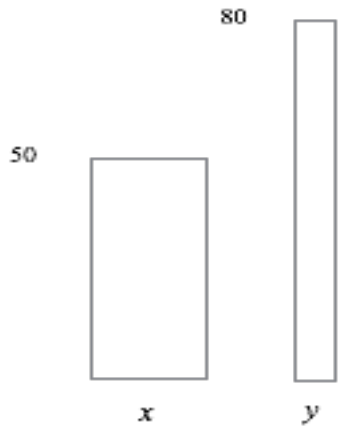

Person *x* will live for 50 years and each of her years (number of years multiplied by wellbeing or quality of life per year) is equal to 1 QALY; while person *y* will live for 80 years but each of his years amounts to only ½ a QALY. Were we then to total up the

QALYs for each person, person *x*'s life will result in 50 QALYs, while person *y*'s life will result in only 40 QALYs. Thus the QALY theorist should choose life *x* over life *y*.

Although I called the QALY method a kind of utilitarianism rendered as economic measurement, QALY is a consequentialist-based theory of resource distribution, but not strictly a utilitarian one, because it chooses between distributive policies based solely on the amount of benefit derived from the consequences of that distribution, whereas for the utilitarian the goal is maximising wellbeing. For the QALY theorist the goal is adding life years, or getting the most cost-efficient life years per healthcare dollar spent. Comparing utilitarianism to the QALY method of distribution, utilitarianism says that we should distribute life extending technology in a way that maximises wellbeing, and the QALY method says we should distribute life extending technology in a way that produces the most healthy life years per dollar spent. Thus the QALY method is a consequentialist view but not a utilitarian view and should not be confused with utilitarianism. QALY theory is a peculiar kind of consequentialism, where the good is not wellbeing as such, but the ideal quotient of healthy years of life that can be bought for the least amount of healthcare dollars. As such, it is primary used as a theory of healthcare economics.

The QALY approach, however, manages to avoid some of the problems with the general utilitarian approach to distribution. Take, for example, the problem of the existence of potential utility monsters, those who could derive an enormous amount of wellbeing from the same resources that would have been allocated to us. The classical utilitarian would have no choice but to distribute resources to these few wellbeing gluttons because this would maximise total wellbeing. The QALY method, however, is based solely on projected quality of life and/or life expectancy and thus it does not matter whether particular individuals, because of their particular preferences, personality qualities, whether they feel more fulfilled in life, etc., would derive more happiness from the same amount of resources given to someone else who has exactly the same life expectancy. The utility monster doesn't have more quality of life, which is a healthcare concept referring to an objective measure of wellbeing that includes many different factors such as the individuals' environment, socioeconomic situation, physical and mental health, etc.

Age is an important determinant of any QALY measurement, because the quality of life and life expectancy (or likelihood of treatment success) are greatly influenced by the age of the individual, with older persons usually scoring lower on

these measurements. Therefore, by distributing medical resources on the basis of the QALY method, priority is given to those with higher life expectancy, other things being equal, because distributing resources to these individuals would maximise the benefit that those resources could potentially bring. The QALY approach, therefore, may come out against trying to extend the lifespan, because this goal is too expensive and yields too small a benefit per resources contributed. Rather, what yields the greatest healthy life year per dollar more often than not might be intervening in the lives on young people in order to improve their health and opportunity for continued life. Older persons, or people with the lowest life expectancies, do not get priority on the QALY distributive theory. Some might argue that this is unfair, that these principles of distributive justice are ageist, i.e. unjustly discriminatory against the old in favour of the young. But the decision depends on empirical contingencies, i.e. which option produces the best cost/benefit outcome. QALY theory could, of course, turn out in favour of life extension too, if the length of projected extension of life and the projected quality of the extended life are ideal.

In the end, the utilitarian and QALY theories are limited as theories of justice because both are equivocal in their answer to how we should distribute the life extension pill, and as such neither approach gets us much further along in fulfilling the demands of justice left open to us by the original libertarian distribution of life extension.

6.6 Egalitarian distribution of life extension

Before I discuss Rawlsian principles of justice and specifically the ‘difference principle’, which I will defend as the most plausible principle for the distribution of resources, I want to examine what a simple egalitarian view says about how we should distribute the life extension pill. This is motivated by the fact that the main objection to life extending technology in the literature seems to be based upon egalitarian intuitions, that, other things being equal, equal distributions are better than unequal ones. Since there are many different versions of the egalitarian doctrine, all giving different weight to the principle of equality, I will focus on a particular kind of egalitarianism, ‘lexical’ egalitarianism, according to which equality is the most important value and given

priority over all other values, i.e. no amount of liberty or utility can justify any loss of equality. This is quite a strong egalitarianism and will be useful in comparing it to a more pluralist egalitarian view, such as Rawls's, in the next section.

Lexical egalitarianism claims that everyone has the right to enjoy the same amount of wellbeing or resources, so if some person has more wellbeing or resources than another person then we should redistribute these good things from the person who has more so that everyone has the same amount of them. Thus lexical egalitarianism implies 'levelling down'. There is something intrinsically good or just about equality, says the egalitarian, thus in order to preserve equality as such they are committed to doing away with inequalities to the extent that we ought to bring down the welfare of persons, e.g. by redistributing their resources, to the level of the worst off person if this is the only way to do away with inequality. For example, say we have two populations, (1) and (2), each containing two people, with the distribution of wellbeing shown below:

Table 6.1 Levelling down

	P1	P2
(1)	110	100
(2)	100	100

(2) is better than (1) according to the lexical egalitarian because (2) contains more equality than (1). Since equality gets moral priority over other values such as utility, then we ought to bring down the wellbeing of the better off person 1 in (1). This move is counterintuitive to most people because the person in (1) enjoying higher level of wellbeing seems like a good thing and to remove some of her wellbeing for the sake of removing all inequalities looks on the face of it unjust.

Consider another example of levelling down which is more sophisticated that can be derived from 'the divided world' case (Parfit 1997). In this example we discover a new continent with people even worse off than the worse off in the known society. We cannot reach this new continent nor can they reach us. Thus the fact that their wellbeing is lower than that in our society has no 'bad effects' on us, and vice versa. But it seems

that the lexical egalitarian must now find this situation unjust and be compelled to level down the welfare in society so that there is global equality between the newly discovered continent and us. The inequality between continents is an impermissible situation for the lexical egalitarian because on this view inequality is intrinsically bad, even if the fact of inequality has no bad effects. As such for the lexical egalitarian, levelling down to equality is always better in some respect, even if we don't have to level down because of the benefit derived from other values, for example, greater utility. This is surely non-intuitive. A more sophisticated pluralist egalitarian, on the other hand, might be able to accept the inequality revealed by the existence of the new population with the lower wellbeing, because even though it still finds the situation now worse in respect to the inequality that exists between the two populations, the fact of inequality existing is not intrinsically bad, thus it is able to avoid the odious commitment to levelling down the welfare of those slightly better off in society because of the good effects of, for example, the rise in the level of total global wellbeing resulting from the existence of the new population.

When it comes to developing a life extension pill, lexical egalitarians, perhaps alone in this, would have an in principle reason not to develop it, since it will not be possible, at least in the beginning, to provide everyone access to it, thus developing a life extension pill would be impermissible. In the likely event that a life extending intervention, such as a pill, is developed, the only just distribution of the life extension pill, according to egalitarianism, is to redistribute it in equal quantity to everyone. If one person is able to have more healthy life years than another person, for example because she can afford a larger amount of the pill, then her liberty must be constrained in this respect and her surplus share of the life extension pill redistributed so that no inequality exists.

It should come as no surprise, then, that lexical egalitarianism is in conflict with the libertarian distribution of resources, that in fact, the egalitarian would maintain that a libertarian distribution of the life extension pill is impermissible. But the lexical egalitarian approach is less intuitively plausible than the libertarian view. I have suggested earlier in the chapter that, while a free market arrangement of the life extension pill is not impermissible, it is not a distribution that is optimally satisfying justice, i.e. it is still an 'unfair' arrangement, and we might be able to find a better theory of distribution. The egalitarian view, at least as I have formulated it, is committed to levelling down and to removing all inequalities even in the face of

substantial gains in wellbeing. A new resource, which is not initially available to all, should be prevented absolutely, a conclusion that I would argue is highly questionable. However, there are more plausible 'pluralist' egalitarian approaches to distributive justice that avoid the problems with lexical egalitarianism and point us towards a fairer arraignment with regard to how we should distribute the life extension pill. I turn now to one such view.

6.7 Rawlsian principles of justice

Although equality should not be the exclusive value defended by a plausible theory of distributive justice, it is still an important value to promote and the one we intuitively equate with justice. John Rawls (1971) revived interest in an egalitarian idea of justice based not on maximising utility or promoting equality as the only value but on 'fairness' in the distribution of basic resources. Importantly, on his view, inequalities in resources should only be tolerated if unequal distribution benefits the least well off members of society. This being said, I want to extend Rawls's theory on the assumption that his view is not really an egalitarian view but rather a prioritarian one, and that the difference principle, Rawls's principle of fair distribution of resources, can be interpreted as similar to the priority view, the view that benefiting those worse off should be a priority. This will become clearer by the end of my discussion of Rawls's theory of justice in this section.

Rawls expresses his central idea of justice as 'fairness' in the following statement: "all social primary goods" (by which he means 'resources', as opposed to 'utility' or 'wellbeing' that the utilitarian approach took to be the values we ought to be promoting), "—liberty and opportunity, income and wealth, and the bases of self-respect—are to be distributed equally unless an unequal distribution of any or all of these goods is to the advantage of the least favoured" (1971: 303; 62). The idea here is not to remove all inequalities (the egalitarianism view, as we saw in the previous section, has highly counterintuitive implications) only those that disadvantage someone. Giving priority to the least well off is about fairness; it is not about treating people equally, because what is just or what is fair is sometimes only brought about by treating people or groups differently, e.g. those with low life expectancy compared with those

with higher life expectancy. This idea will come up again when I discuss Norman Daniels's view of justice between the young and the old in Chapter 8.

The basic intuition behind the Rawlsian idea about justice was that justice is not satisfied by simple equality of opportunity, which is the idea that people should not be discriminated against on the basis of race, gender, or disability in their access to basic social goods. Something further is needed. This is where the idea of fairness comes in, that the least favoured in society have certain rights and freedoms that must be protected and given priority in the distribution of opportunity and resources. As a result, justice, according to the Rawlsian, is 'fair equality of opportunity'.

The initial defence Rawls gives in favour of his idea of justice as fairness is to appeal to the intuition that the differences between us, e.g. intelligence, natural talent, are morally arbitrary and that justice should not care about them; they are not the result of a persons' choices or work, but merely a matter of luck. The more naturally talented do not deserve their advantages if they do not in some way compensate the least well off, just as those who are not naturally gifted do not deserve their disadvantages. It is distributing basic resources (not just economic resources, but other basic goods that everybody needs) fairly that matters morally, which is maximising the share of the person at the bottom.

The second defence that Rawls gives of his conception of justice as fair equality of opportunity comes in the form of an adaptation of the contract argument. He modifies the social contract conception of justice and uses it as a hypothetical device to tease out intuitions people have about how to order a society and social institutions.[51] He calls his social contract the 'original position', which can be described as an artificial 'state of nature' where all individuals are obscured behind a hypothetical 'veil of ignorance' so that none may know anything about their particular circumstances, e.g. whether one will be male or female, endowed with particular natural talents, have particular interests, etc., anything that allows those born with such natural advantages to further establish those natural advantages. From such an 'original position' people have to design a set of rules or principles for society.

These principles of justice will be what any rational person would agree on if they were innocent of their particular circumstances and thus reflect what is universally fair, says Rawls. They are meant to be 'universalisable' principles of how society and

[51] Traditional conceptions of the social contract derive from the work of Hobbes *Leviathan*, Locke *Two Treatises of Government*, and Rousseau *The Social Contract*.

social institutions should be set up that any rational person would choose, principles of justice you would will from any position or circumstance, regardless of the consequences.[52] The rational strategy that people are forced to use in this situation is an economic principle not original to Rawls called 'maximin', or maximising the minimum share. As we will see below the maximin rule leads to Rawls's principle of distributive justice, the 'difference principle'. Rawls thought the rules of society designed by people in the original position would constitute these two principles (1971: 302-303):

> *First Principle*: Each person is to have an equal right to the most extensive total system of equal basic liberties compatible with a similar system of liberty for all; *Second Principle*: Social and economic inequalities are to be arranged so that they are both: (a) to the greatest benefit of the least advantaged, (b) attached to offices and positions open to all under conditions of fair equality of opportunity.

The first principle is simply the claim that people should have their basic liberties protected, which is usually not contested by other theories of justice and is a principle that Rawls's theory shares with libertarianism, but libertarianism rejects redistribution in order to rectify inequalities of circumstances or resources. The second principle of justice chosen within the original position, which is divided into two sub principles (a) and (b), is more contentious. 2(a) concerns the distribution of resources, and 2(b) the government of the distribution of opportunities. Rawls puts 2(b), the value of fair equality of opportunity, in a higher lexical priority to 2(a), the value that economic inequalities ought to be redistributed so as to produce the greatest advantage to the least favoured person. Rawls calls the principle that this distributive strategy leads to the 'difference principle'. The basic idea behind the difference principle 2(a) is the same as Rawls's general conception of justice stated earlier, that all primary social goods should be distributed equally unless an unequal distribution is to the advantage of the least well off. This basic principle requiring us to focus on the condition of the worst off in society is the centre of the Rawls's theory of justice. It is the principle I will use as representative of the Rawlsian approach to the just distribution of resources on the issue of how to fairly distribute life extension that follows in the next section.

Rawls claims that the difference principle aligns with our intuitions that justice in the distribution of resources is found from the standpoint of the most adversely affected person, the least advantaged in society. If I were the least well off member of

[52] Rawls has called his theory of justice a 'proceduralised' interpretation of Kant's categorical imperative, 1971: 226.

society and I would still choose a particular principle of distribution of basic resources, this indicates that we are getting at some universalisable principle of justice and perhaps a fairer alternative to the principle of maximising utility or any other basic principle of justice. Thus the standpoint of justice, according to Rawls, is the standpoint of the least well-off person in society, not equality per se. In this sense, the question of justice is actually a self-interested question you ask when in the original position: what system would I pick so that I am in, at least, a not disadvantaged position when I later discover what my place in society is, my socioeconomic status, natural talents, preferences, and so on. So no matter what your position in society is, what preferences you have in life, you want more resources rather than less resources, more liberty rather than less liberty, etc., and so you will prefer the principle that guarantees this. Thus Rawls's arguments are arguments for this central principle of justice, which obviously does not map onto a pure egalitarian view of justice as such.

We must try to understand the lexical ranking Rawls applies to his principles of justice. Since his principles of justice affirm more than one value, e.g. maximising liberties, and fairly distributing income and wealth, he needs to say which values trump other values when they come into conflict over policy decisions. Rawls puts the value of ensuring people have "equal right to the most extensive total system of equal basic liberties" first in his lexical priorities, second he puts the value of fair equality of opportunity to social and economic resources, and third is that, when it comes to social and economic resources, we should maximise the share of the least favoured members of society. In regard to Rawls's first principle, there is no reason to believe that the development of life extension would threaten peoples' basic liberties, thus I will assume for the remainder of the discussion that this principle is satisfied in this context. The difference principle, which is actually third in Rawls's ranking of his principles of justice, is the principle of justice that, I suggest, is most relevant in discussion of how to best distribute a life extension pill.

That being said, there is at least one major objection to the difference principle as a theory of how we should distribute resources. The difference principle, it might be objected, cannot tell us which of two possible distributive outcomes is better if both contain the same level of wellbeing for the least well off. For example, how would the difference principle help us decide between two outcomes with different distributions, both equal in the wellbeing of those least advantaged while unequal with regard to the level of wellbeing of the most well off?

Table 6.2 The difference principle and choosing between outcomes

d	3	5
c	3	4
b	2	2
a	1	1
	A	**B**

Table 6.2 shows two outcomes A and B, each with four people ‘a’ through to ‘d’. The numbers represent their level of wellbeing, from the least advantaged at the bottom to the most advantaged at the top. In both outcomes A and B the least well off have the same wellbeing, these are persons ‘a’ and ‘b’ in ascending order of worse off. However, in outcome B, the two most well off persons enjoy a higher level of wellbeing than the two most well off persons in outcome A.

The first thing to say is that it might look like the difference principle cannot advocate one distribution over another, because the least well off share the same amount of wellbeing in both outcomes, even though one outcome, outcome B, clearly has the greater total wellbeing. But there is a way, using the difference principle, that one might start to decide between outcomes in the example illustrated above. What the Rawlsian does is apply the difference principle all the way up from the least advantaged until they reach the most advantaged, by moving on to the next least advantaged and compare their level of welfare, then the next, and so on. However, by scaling up the difference principle we only get so far in the above case. Person ‘c’ in outcome A becomes the least advantaged if ‘a’ and ‘b’ (whose wellbeing is equal in both worlds) didn’t exist. And since ‘c’ and ‘d’ in A are ‘equally the least well off’—thus equality has been fulfilled among the least well off—then it seems the difference principle would tell us to choose outcome A over outcome B. Because if persons ‘a’ and ‘b’ didn’t exist in outcome B, the least advantaged person would be person ‘c’ with a level of wellbeing higher than the level of wellbeing of the least advantaged (person ‘c’) in outcome A.

This is where my interpretation of the difference principle departs from Rawls. If healthy life years are considered good, then our aim should be to maximise the healthy life years of the least advantaged person in society. Thus, on my version of the difference principle, we should choose outcome B over outcome A. This conclusion, I suggest, is more attuned to our intuitions in the case above.

6.8 The difference principle and the distribution of life extension

We have seen how the Rawlsian difference principle is less problematic than libertarian, utilitarian, and lexical egalitarian distributive principles and can now turn to applying this principle to the question of how we should distribute the life extension pill.

If healthy life years is considered good, then our objective should be to maximise the healthy life years of the least advantaged person in society, but the least advantaged persons in this respect are most likely to be those in society who do not have their basic healthcare needs taken care of, who cannot expect to reach even the average life expectancy let alone have to worry about extending their maximum healthy lifespan. Consequently, when it comes to the question of whether we should fund life extending interventions, Rawlsian principles of justice gives this a low priority.

So, if we were guided by the difference principle we would not give priority to developing life extension because there is little benefit to those least advantaged in society in doing so. Nevertheless, developing the technology would have significant benefit for better off members of society. If a life extension pill were to be developed, we can ask how the difference principle would best distribute such a resource. At the start of the chapter I told the story of how the life extension pill would come into the world: the pill was developed privately, thus a free market arraignment would obtain at least in the beginning, so those with the money can freely purchase the life extension pill. Eventually, seeing the benefit of a wider distribution of the life extension pill, the government steps in and buys up the entire supply of the life extension drug. We return now to the story, in which the situation is that the government has the life extension pill, but of course there is not enough to give to everyone who might want to use the drug, and so must now decide which individuals or groups to give priority to for the drug. The goal of distributive justice is to distribute the life extension pill equally unless an

unequal distribution makes the least advantaged person as well off as possible. As the life extension pill increases life expectancy, then we might give priority for life extension to the most disadvantaged groups in society with respect to life expectancy. But the most disadvantaged groups in society in respect to life expectancy are those who die earlier from diseases and conditions related to lower socio-economic environments. As such a life extension pill will not have much effective on their life expectancy. But for those not otherwise afflicted by disease, the life extension pill will benefit some more than others in respect to life expectancy, and some of these are older than other, i.e. better off with respect to the goods of life. Therefore, my suggestion to the question of what justice says about how the government should distribute the life extension pill is that: the life extension pill should be distributed (among those groups to whom it will be effective) equally unless giving priority to those least advantaged with respect to life expectancy (those less better off with respect to the goods of life—more discussion of this in Chapter 8) makes them better off. Consider some policies.

Fig. 6.1 Distributing the life extension pill

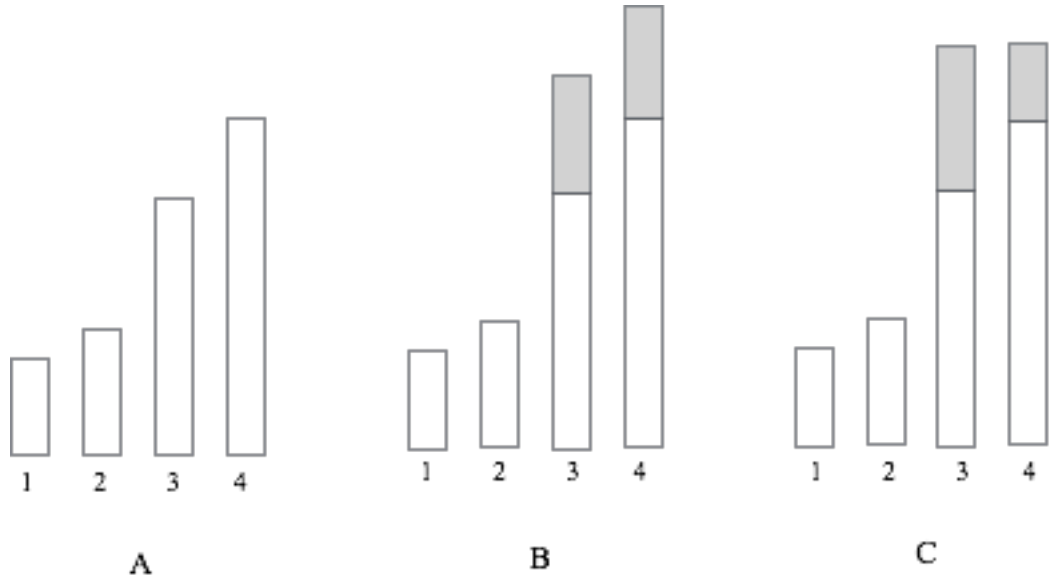

Figure 6.1 above depicts three separate policies the government has devised for the distribution of the life extension pill. Numbers 1 through 4 represent people of different advantage in society with respect to life expectancy, while the height of the blocks represents projected life expectancy. The greyed-out areas extending above the blocks is the increased life expectancy provided by the life extension drug. In all policies A, B,

and C persons I and 2 die early of disease and the life extension pill cannot help them. In policy A, the government chooses to hold on to the life extension pill and not distribute it. In B those on whom the pill will be effective (persons 3 and 4) are given equal doeses of the life extension drug. And in C, priority is given to person 3 (the least well off in the situation, when 1 and 2 are excluded) who is given a maximal dose of the life extension drug, the most that can help, and person 4 (who is already highly well off in regard to life expectancy) is given the remainder of the drug, which, in the end, makes them equal in life expectancy with person 3.

The lexical egalitarian, given these three policies, would have to prefer policy A, since there is less inequality in A than in either B or C. The utilitarian would advocate policy B, because there is greater total wellbeing in B than in A or C. And lastly, a government guided by my prioritarian interpretation of the Rawlsian difference principle would find policy C the best option, as it gives priority to person 3 (the worst off among those the life extension pill will benefit) and brings their life expectancy up to the level of the most well off in regard to life expectancy, the latter still able to get some of the life extending drug. Person 3 is better off in policy C than in policy B, while person 4 is worse off in C than in B, but policy C is the fairest policy out the three. (Persons 1 and 2 are not made worse off in either, except in the fact that the gap between their life expectancy and those of 3 and 4 is greater).

Those in the least advantaged strata of society will likely have little use for the life extension pill if their basic health needs have not been satisfied first. On the other hand, the rest of society, who already have a chance of achieving high life expectancy, might well derive great benefit from augmenting this by using the life extension pill. Moreover, it would be a hard sell to argue that the most well off are disproportionately made worse off because they did not get priority for the drug and thus received less of a share of life extension than they could have had. Under my interpretation of the difference principle, so as long as, out of those whom the life extension pill can significantly benefit, those least advantaged as to life expectancy get as good a share of life extension as possible, then fairness has been achieved.

6.9 Conclusions

I have argued that, while inequalities in the distribution of life extension resulting from a libertarian free exchange are not inherently unjust, justice might be fulfilled more optimally through redistribution based on a further principle of justice, this was a modified Rawlsian difference principle. I disagreed with the libertarian by claiming that a non-libertarian redistribution can be just if it makes things better or fairer. This non-libertarian redistributive principle, I concluded, is a prioritarian interpretation of the Rawlsian difference principle, which says that, among those whom life extension will benefit, priority for life extension should be given to those least advantaged in regard to life expectancy.

The result of all this is that, not only is developing life extension not impermissible, nor unequal access to it intrinsically unjust, but that there might be an optimal redistribution of life extending technology (once it has been developed) in the form of giving those with the lowest projected life expectancy priority access to it, while those more well off in this respect are not prohibited from accessing the drug.

CHAPTER 7

NONCONSEQUENTIALISM AND LIFE EXTENSION

> Human rights do not get any more fundamental than the right to carry on living.
>
> Aubrey de Grey, 'Life extension, human rights, and the rational refinement of repugnance' 2005: 5

My argument for the permissibility of developing life extension in Chapter 5 was a consequentialist argument. But I am not defending a consequentialist view; rather I am surveying what different moral theories have to say about life extension. We saw in Chapter 6 different ways we might think about what would count as a just method of allocating life extension, some of which were non-consequents distributions. Might the permissibility of life extension be further argued for from non-consequentialist grounds? For example, could we simply have a duty to provide people with life extension if such was their wish? This is the question I wish to briefly explore in this chapter. I will proceed by using a general Kantian approach to the question of whether we ought to develop life extension and thereby allow people to prolong their lifespans, ultimately concluding that full and autonomous persons have a 'conditional right' to life extending technology.

7.1 Do we have a positive right to life?

In this section I will argue that persons have a conditional right to prolong their lifespans. But what is a right to life? A positive right to life is essentially an obligation on others that one's continued life be promoted. It can be distinguished from a negative right to life, which can be understood as a right not have one's life ended (to be killed) by some other agent. Clearly the negative right to one's life is less demanding than the former positive right to have one's life prolonged or promoted in some way. I ultimately want to argue that so long as a person possesses a right to life in the positive sense, they also have a right to the resources to prolong their lifespan, other things being equal.

My first task is to establish whether persons have a positive right to life. I will do this by first distinguishing the value of the life of a 'person' from the life of non-personal beings. This will be accompanied by discussion of utilitarian and Kantian accounts of the value of the life of a person. I will defend a broadly Kantian view, especially relying on a normative principle of respect for autonomy and suggest that persons can be possessors of a positive right to life based largely on a duty to respect the desires of autonomous persons. In section 6.2 I will finish with the qualification that, while I acknowledge that the positive right to life of persons entitles them to life extending interventions, this does not necessarily commit me to the thesis that we have a moral obligation to extend the lifespan in all cases. We have, rather, a prima facie duty to extend the lifespan, other things being equal. This is because the positive right to life is conditional largely on the desires of the person to continue living. Thus the conflicting rights of other persons, and of other considerations, regulate the positive right of any person to go on living.

7.1.1 The value of the life of a person

We often distinguish morally between the life of a rational being and other kinds of life by claiming that killing or allowing a rational being to die would be worse than killing or allowing some other kind of being to die. What accounts for the special value we attach to the survival of a thinking or reflecting being? The answer lies in our intuitive conception of a thinking being, which I call a 'person', and the valuable features inherent only in the lives of persons.

We get from John Locke the traditional analysis of the notion of a 'person'.[53] Locke proceeds by distinguishing between non-human and human animals. 'Human being' refers to that aspect of the organism which is a living body of a particular biological kind, i.e. a human is a kind of animal, a sentient animal, but not necessarily a person. A person, however, is a particular kind of intelligent being, possessing rational qualities, and most importantly, an understanding of itself as existing over time. Persons are not necessarily human and humans are not necessarily persons.

[53] Locke 1694, *Of identity and diversity*, in Perry 2008.

Unlike the human animal, which is only related to past and future beings by sameness of 'life' or biological body alone, the sameness of a person is to identity of 'consciousness'. A person, or 'personal being', is "a thinking intelligent being, that has reason and reflection, and can consider itself as itself, the same thinking thing, in different times and places" (2008: 39). We can see that persons are capable of more than mere sentience, the experience of pleasure and pain, but also self-awareness, a capacity to reflect on one's past and future, and in particular the capacity to have desires in regard to one's future existence. We might reasonably infer that non-human animals do not typically share these features of reflective agency and that certain kinds of human beings, such as new-borns, infants or severely mentally disabled human beings are in the same category. Therefore, not all human beings are persons, and conversely, not all persons need necessarily be human beings in order to possess the above characteristics of a rational, self-aware being.

Locke's general conceptualisation of a person also carries with it normative implications. The key difference between a person and a non-personal being, as described by Locke, is the possession of 'consciousness'. Locke did not mean consciousness in the 'weak' sense of being merely sentient, for even non-persons have sentience or are 'conscious' in the weak sense. He means that personal beings are conscious in the stronger sense of being 'conscious to oneself', of having knowledge of oneself, knowledge of one's own past thoughts and actions, one's desires and intentions, a kind of reflective agency. He says: it is "the sameness of a rational being; and as far as this consciousness can be extended backwards to any past action or thought, so far reaches the identity of that person; it is the same self now as it was then" (2008: 39-40).

Beings that are not persons do not possess this self-reflective awareness of themselves as existing over time, that 'my' past actions are connected to 'my' present ones and these to 'my' desires about the future. Consequently, only a rational being who possesses this 'consciousness of oneself and one's actions' can possibly be held responsible for their actions. Thus we tend to include persons in a special normative sphere of consideration that would be inappropriate in regard to other kinds of beings. As we shall see later, this is also why a person is a candidate for a particular kind of positive right to life, or a right to have a say in their own continued existence. Of course, I am not arguing that persons are the only beings who have rights, rather that

they are the only kinds of being who satisfy as candidates for the particular positive right to be kept alive.

7.1.2 Different moral views on the value of the life of a person

The normative implications attached to the life of a person on the Lockean account is similar to what Kantians find valuable in the life of a person, why a human being, in having a rational nature, is uniquely special. For the Kantian it is the rational nature in human beings that is an end in itself. Kant says, "rational nature exists as an end in itself", and "the human being necessarily represents his own existence in this way" (*Groundwork* 4:429).

The Kantian notion of a rational being possessing a self-reflective nature, what we've been calling a 'person', is a being that can think of her own existence as an end in itself. We saw hints of this in the Lockean account of a person being conscious of her own past actions and an awareness of how these affect and relate to one's continued existence. For the Kantian, fundamental moral status is granted to persons because of their possession of a rational nature, this ability to be autonomous and 'self-legislating'. This valuing of the rational nature in the existence of human beings, or 'persons', led Kant to formulate a principle regarding our fundamental moral duty towards such beings, his Formula of Humanity: "So act that you use humanity, whether in your own person or in the person of any other, always at the same time as an end, never merely as a means" (4:429). The reasoning behind this principle being that, if rational nature is something worth valuing intrinsically as an end in itself, then beings possessing a rational nature, i.e. persons, ought to be respected as ends in themselves.

Ultimately, I want to defend a Kantian conception of the value of the life of a person, and in doing so, show how this conception lends weights to the argument that persons possess a positive right to go on living. But first I will look at what the consequentialist has to say about the value of the life of a person and attempt to show some of the shortcomings of the popular preference utilitarian view in this regard.

We have seen that, in particular, the life of a person has greater moral implications than the lives of other kinds of beings. For instance, we can find in the characteristics that identify the kinds of beings who qualify as persons indicators of why it would be worse to kill a person than to kill a non-person. The Kantian can argue that

it is wrong to end the life of a person because a rational nature is worth preserving in and of itself. Therefore, the continued life of a person is morally valuable in this sense so long as it continues to be the unique possessor of features like autonomy and self-awareness, that is, continues to be the life of a person. Consequentialist views agree with the Kantian view in being able to recognise that the life of a person is different from that of a non-person and seems to warrant greater moral consideration, but differ from the Kantian view in the way they value what is special about the life of a person.

Consequentialist moral theories say our only moral obligation is to maximise the good, which in the case of classical utilitarianism, is pleasure.[54] But the utilitarian theory has gone through changes in regard to what is 'good in itself', the good to be maximised. For classical utilitarians it was pleasure, for Bentham this meant any basic hedonistic kind of pleasure, for John Stuart Mill it meant somewhat 'higher' pleasures or mental states, the kind of intellectual pleasures open to a Socrates but not to a sentient animal.[55] Later utilitarians, such as Henry Sidgwick thought that the satisfaction of a person's preferences was the intrinsically good thing to be maximising, and modern preferences utilitarians, like Peter Singer, adhere to this idea of the good.[56]

Of course, not only persons experience pleasure and pain, but the preference utilitarian recognises that there are certain kinds of pleasure that only a person can experience, through those features of personhood like rationality, reflection and self-awareness and the capacity to form desires. But these capacities are not valued in and of themselves, their value for the preference utilitarian is instrumental, i.e. that personal beings are able to generate more sophisticated kinds of value than pleasure, namely, they have the capacity to form preferences about their own conception of the good life, and it is a person's preferences that we must aim to satisfy. And this is the path the preference utilitarian takes to the view that it would be worse to kill a person than to kill a non-person.

Consider the following. The classical utilitarian is able to see that the continuation of a person's life beyond the present moment holds for them the expectation of a certain amount of pleasure. Since pleasure is based on experience only, if the person dies suddenly and painlessly, even with the prospect of a future and preferences regarding their own future existence, no badness attaches to this loss of

[54] Bentham, Jeremy. [1789]. 2000. *An Introduction to the Principles of Morals and Legislation*. Batoche Books.
[55] Mill, J.S. [1863] 1957. *Utilitarianism*. 2nd (rev.) ed. New York: Liberal Arts Press.
[56] Sidgwick, H. *Methods of Ethics* [1874]; Singer, P. *Practical Ethics*. 1993.

potential good, other things being equal. So the future prospect of a person's life offers the classical utilitarian no direct reason to value the continued life of the person, no direct reason against killing the person. However, a preference utilitarian holds a direct reason for valuing the continued life of a person based on the satisfaction of the future-orientated desires of a person. All else equal, you shouldn't kill people, if this will lead to less net pleasure in the world. But on classical utilitarianism, people are 'replaceable', thus it is permissible to kill them if their deaths will lead to the creation of similar beings who, in their lives, will 'make up' for the lost pleasure of those who died. However, on a preference utilitarian view such as Singer's, people are not replaceable.

Not only that, but to have their life prematurely terminated would not satisfy the person's future-oriented preference to continue living, i.e. replacement, bringing into existence a new person, won't compensate for the death of a person. In this way a preference utilitarian can view the killing of a person as worse than the killing of a being that has no capacity to form a desire to avoid the future termination of its own existence. We can see this view expressed in the following passage from Peter Singer's defence of preference utilitarianism (1993: 95):

> For preference utilitarians, taking the life of a person will normally be worse than taking the life of some other being, since persons are highly future-orientated in their preferences. To kill a person is therefore, normally, to violate not just one, but a wide range of the most central and significant preferences a being can have...In contrast beings who cannot see themselves as entities with a future cannot have any preferences about their own future existence...such beings might struggle against a situation in which their lives are in danger...but this indicates no more than a preference for the cessation of a state of affairs that is perceived as painful or frightening.

Preference utilitarianism, in saying that the good we must promote is the satisfaction of people's preferences, is, in this sense, able to accommodate a person's own conception of the good in their own lives. Though the right action is still the one that maximises the good under both versions of utilitarianism, it is the conception of the good that is different, and by making the good to be maximised the satisfaction of a person's preferences, the preference utilitarian is able to avoid certain shortcomings of the classical view which says the only good to be promoted is pleasure.

Nonetheless, we can see that for the consequentialist, the moral distinction between the life of a person from that of a non-personal being merely changes the kind of thing that the consequentialist says we ought to promote. They do not value those special features of personhood such as autonomy and self-awareness in and of

themselves as morally valuable. The preference utilitarian only values the promotion of the good, whether that is pleasure or preference-satisfaction, but she has difficulty in valuing the person, who has the capacity to determine their own conception of the good, as an end it themselves. An example might help here. Consider the standard scenario in which you can stop a trolley from killing five people by pushing a fat man onto the track but killing him in the process. The preference utilitarian, must see it as permissible to kill the fat man and save the other five people, the non-utilitarian on the other hand, is constrained by a respect for persons in and of themselves, thus forbidding them to kill the fat man even in this action would save five other people.

The Kantian, unlike the consequentialist, takes respect for autonomy as an independent moral principle when valuing the life of a person. Respect for persons might also be thought of as acting as a 'side-constraint' on utilitarian action, forbidding us, in some cases, from maximising good consequences (Nozick 1974). Respecting rationality and autonomy, the central normative features of the life of a person, is to respect their intrinsic value as a person. Violating an autonomous person's decision not to be killed, therefore, would be violating the value of the person qua 'person'. From a non-consequentialist approach, valuing the autonomy of a person is not about promoting this feature of a person's life, but rather to respect it in and of itself, respect it as an end in itself, rather than only when it is useful in promoting a certain outcome. Kantians don't 'promote' autonomy, they 'respect' autonomy, just as consequentialists 'promote' the good, they do not 'respect' the good (Pettit 1993).

Though respect for autonomy is largely a non-consequentialist value, autonomy, as such, has value for the consequentialist only in so far as persons have the preference that they continue to be an autonomous being. For example, the preference utilitarian would find nothing wrong with the desire of a person to undergo an invasive medical procedure that would disable his frontal lobe functions in order to regress back to the simple pleasure state of non-self-awareness of a lesser animal; likewise the preference of the person who, upon being newly freed from the Matrix, desired to go back to his blissful existence as a brain in a vat (*The Matrix* 1999, The Wachowski Brothers). The preference utilitarian cannot value autonomy as something special in and of itself, but only so far as valuing autonomy brings about the outcome with the maximum preference satisfaction.

Our normal intuitions that people would act wrongly in cases like the person voluntarily undergoing a lobotomy, or that of one deciding to be plugged back into the

Matrix illusion, would run counter to the consequentialist interpretation. We usually think there is something deeply wrong about wanting to give up our autonomy or higher rational abilities, that the person's decision in both of the above cases somehow violates some intrinsically valuable feature of being a person. Our intuitions suggest that one's autonomy, because it is something worthy of respect in its own right is, in a sense, not one's to violate. These examples, therefore, suggest that a Kantian account of the value of the life of a person may have more intuitive appeal or that it strikes more at what we want to say is important or special about the life of a person.

The above examples can also be inverted. Consider a case where respecting a person's autonomy reduces her level of preference satisfaction: An elderly person decides now, in the present, that she wants to go on living because living the life of a person is a valuable thing. But we know, or can be fairly certain, that any continued life in her future would, in fact, be degrading to her because her faculties of reason and bodily functions would rapidly decline. In this case, continued life would not, in fact, satisfy her preference to live the life of a full person.

Therefore, consequentialism entails that we must be paternalistic, that we must deny her continued life. But in doing so, we violate her autonomy, based on the principle of the respect for the autonomous decisions of a person regardless of the eventual outcome. From the principle of respect for autonomy, it might seem that no amount of reduced preference-satisfaction in the persons' future justifies violating the decision of the, at the time, autonomous person to go on living by prematurely killing her or letting her die.

Again, on the intuitive level, we feel that disrespecting someone's autonomy would be tantamount to violating the value of the person as a person. Accordingly we might find ourselves attracted to a moral theory that places independent value on respect for autonomy in regard to the value of a person's life and the wrongness of killing a person.

Nonetheless, both consequentialist and Kantian moral approaches are able to recognise and put weight on the fact that the life of a person can be seen as something worth living by the person themselves. Only a person can form the rational desire 'I desire to go on living' which would be frustrated were they to be killed. This capacity will form the basis of persons having a right to life in the positive sense.

7.1.3 A person possesses a positive right to life on a Kantian view

Recall Kant's notion of a human being whose rational nature it is to represent its own existence as an end in itself, to be autonomous or 'self-legislating'. We can interpret the Kantian notion of a person representing her own existence as an end in itself as the notion that a person is able to construct what we might call a 'practical identity' over time by the fact that she has knowledge of, or is aware of, her own past actions and thoughts and thinks about her own future. This is an indicator of how far the identity of the person is able to reach, of how their past, and especially their future, holds great importance for her in a way that is impossible for non-personal beings. This unique ability to construct a future-orientated practical identity is relevant to making the case that persons, by their rational nature, may possess a right to continue living.

From a Kantian perspective, in virtue of respecting the person's rational nature in and of itself, a person is both the object of a right to life and a possessor of this right: as an autonomous being she is the object of respect for autonomous existence and she possesses a right for her decision as it affects her existence, her continued life, to be respected. The right to life, then, regarded as something that is 'possessed' by self-aware, autonomous beings, carries the idea that it is something that may be exercised or discarded by them if they should so choose. The right to life is conferred to persons in respect of their existence as autonomous beings who may want to go on living if they desire: a right to be given life; to receive more life. Thus the Kantian, from a principle of respect for autonomy, can account for why persons have a fundamental and independent right to continue living, not merely a negative right not to be killed.

Though the preference utilitarian view can perhaps only say why it would be worse to kill a person than a non-person, namely, a person's preference to continue living would be equally frustrated by being killed as by not being saved. So the preference utilitarian can account for why it would be worse not to prolong someone's life if they desired it. But the preference utilitarian does not go further by turning this into the claim that personal beings must then have a positive 'right to life'.

I would argue that the positive right to life of a person, even in the Kantian sense, is probably not an absolute right, but rather it is a conditional, or prima facie, right that can clash with other rights and other considerations. One reason why a person's right to go on living is a conditional rather categorical one is that, since his value as a person is based on the value of autonomy of being a human agent, a value

which is collectively possessed by other human agents, his value is in virtue of being a member of the group 'moral agents' and so is a value larger than any one person. Kant says something that points towards this idea in a continuation of the same passage from the Groundwork quoted earlier, in which he states the principle that a rational nature exists as an end in itself is the grounds for formulating a universal practical law, i.e. the Formula for Humanity (4:428-4:429):

> The ground for this principle is: rational nature exists as an end in itself. The human being necessarily represents his own existence in this way…But every other rational being also represents his existence in this way consequent on just the same rational ground that also holds for me, thus it is at the same time an objective principle from which, as a supreme ground, it must be possible to derive all laws of the will. The practical imperative will therefore be the following: So act that you use humanity, whether in your own person or in the person of any other, always at the same time as an end, never merely as a means.

Consequently the right to life of an individual person may sometimes be overridden by the right to life of other persons based on their own autonomous nature in certain circumstances. What all this amounts to, it seems, is that the rational nature of a person upon which the right to life is based is merely a feature that gives us a moral reason to support a claim to a right to life in certain circumstances, but makes it a right that is conditional on the similar right of other persons.

7.2 A conditional right to life extending technology

Guided by a Kantian principle of respect for autonomy, one of the primary ways to respect a person's autonomy is to respect her desire to go on living, which can be done by entitling her to a right to life in the positive sense. We see this when we apply the right to life in the context of the older persons' desire to go on living. So long as the older person is in possession of a right to life, in the non-consequentialist sense of an entitlement granted to autonomous agents, this right to life should in the very least allow the possessor to choose to continue living as an autonomous agent.

Moreover, if there exists a right for the young to the medical technology to prolong their lifespan, why not the same right for the old, irrespective of years already lived? If both desire equally to go on living, and both equally have a positive right to life, then from a principle of respect for the autonomous decisions of persons, we must

provide the resources to facilitate that right to life to both equally, other things being equal. From the perspective of healthcare, life-prolonging medicine effectively becomes medicine for the old, just as lifesaving medical treatments for the young person are, in a sense, life-prolonging treatments. In this sense there is no difference between life extending interventions for the young and life extending interventions for the old. So long as the old still autonomously submit their right to go on living, then we fulfil our duty towards them in regard to this right to life when we provide them with the medical technology to prolong their lifespan, again, other things being equal.

The right to life includes a respect for one's decisions as they bear on one's life. If this is true, then in order to respect a persons' autonomous decision to prolong her life we must provide her with the means to do so, i.e. the medical and technological resources. This is what a positive right to life demands. However, as I have suggested above, the positive right of an individual person to continue living may sometimes be overridden by the right of another person to continue living in certain circumstances based on their own autonomous nature. Thus the right to the means to prolong one's life is also conditional because of the competing right to life extension of other persons. This conclusion might not seem like it gets us very far, especially in regard to the question of who should get priority for life extension when resources are limited and a persons' conditional right to prolong their lifespan is in competition with a similar right of others.

A final word regarding the implications of a person's positive right to interventions to prolong their lives. I would claim that we are talking about two different kinds of obligation when we talk about a right to the technology to prolong the lifespan, and a moral obligation to prolong the lifespan. The former is a conditional right that those entitled may exercise or discard as they choose, so long as this does not clash with similar positive right to life of other persons under particular moral approaches. The latter implies an obligation to always prolong the lifespan regardless of individual desires, in a similar vein to those who argue that 'life is sacred' and that we are always obligated to keep people alive, a claim I would reject. Therefore, the 'right' to prolong one's life is conditional on, one being an autonomous agent capable of having a preference for more life, and the prolongation of one's life (including the access to the technology to do so) does not negatively encroach on the conditional right of another person to prolong their life.

CHAPTER 8

LIFE EXTENSION AND A 'FAIR INNINGS'

> ...to treat the older person, letting the younger person die, would...be inherently inequitable in terms of years of life lived.
>
> Michael Lockwood, 'Quality of life and resource allocation' 1988

The desire to avoid death by extending our lifespan is not such a pressing desire for young persons because they, generally, have many years of healthy lifespan in front of them and are not yet knocking on death's door. For older persons, however, the desire is much more acute, at least for those who have strong categorical desires compelling them into the future. Given that the desire and need for life extending interventions is generally stronger for older persons than for younger persons, do the former deserve priority when it comes to access to life extension?

I argued in Chapter 6 that governments distribute the life extension pill fairly by giving priority to the least advantaged with respect to life expectancy among those whom life extension will be effective. This could sometimes mean that older persons get priority over younger persons. However, the dominant view in the context of healthcare distribution (though not restricted to this context), is supported by the strong intuition that, when in competition, life extending interventions should go to younger persons in favour of older persons because older persons have already had a fair share of life. The 'fair innings' claim says that, once a person has lived a certain number of years (has had a 'fair innings') they have less right to continued life than those who have not yet reached a fair innings.[57]

Thus, according to the fair innings view, priority for life extending interventions ought to be decided based, not on life expectancy, but on age. Unlike the prioritarian interpretation of the Rawlsian view of distributive justice that I argued for in the previous chapter, the fair innings view is founded on the idea of life years already lived, disregarding the moral relevance of life expectancy. Both the view I have defended and the fair innings view argue for the 'fairness' inherent in their approaches to the

[57] For discussion of 'fair innings' see: Harris 1985; Rivlin 2000; Williams 1997.

distribution of resources. But which approach is really the fairest? This is the question I will try to answer in this final chapter.

8.1 The fair innings claim

The fair innings claim has been clearly articulated by John Harris (1985: 91):

> The fair innings argument takes the view that there is some span of years that we consider a reasonable life, a fair innings. Let's say that a fair share of life is the traditional three score and ten, seventy years. Anyone who does not reach 70 suffers…they have missed out on a reasonable share of life; they have been short-changed. Those, however, who do make 70 suffer no such injustice, they have not lost out but rather must consider any additional years a sort of bonus beyond that which could reasonably be hoped for. The fair innings argument requires that everyone be given an equal chance to have a fair innings, to reach the appropriate threshold but, upon having reached it, they have received their entitlement. The rest of their life is the sort of bonus which may be cancelled when this is necessary to help others reach the threshold.

I must admit that I find the fair innings intuition very plausible. Because an older person has had or is closer to having had a 'fair share' of life, it is justifiable, given limited medical resources, to favour the younger person. In other words, it is less wrong to trade off the continued life of an older person in favour of the continued life of a younger person because the former have already reached a certain threshold of life demanded by justice and fairness. The question we must ask, then, is whether it is really true that having already lived a certain number of years necessarily makes one less entitled to extend one's life.

The fair innings claim and its prominence in decisions about healthcare distribution has often been used to justify valuing the continued life of an older person less than the continued life of a younger person. Valuing the continued life of an older person less than the continued life of a younger person has been referred to in the literature as 'ageism', not perhaps wholly as a discriminatory view but as a reasoned one. The latter is the sense in which I will use the view 'ageism' in this chapter. I will refer to an ageist view as a view that says that, when saving life, we should give priority to the younger person, not to say that this view is 'ageist' in the sense of unjustifiable discrimination against older persons. However, with life expectancies increasing in developed countries and the possibility of the extension of the lifespan, it could be

argued that ageist intuitions in healthcare distribution are starting to become harmful. If it is possible for older persons to radically extend their healthy lifespan then being older does not necessarily mean less opportunity for continued healthy life. Moreover, the notion of a traditional fair innings becomes problematic.

I will suggest in this chapter that there are two main arguments used to justify valuing the continued life of an older person less than the continued life of a younger person: one of them being the 'fair innings' claim, and the other will be the 'deprivation account' of the badness of death, with which we are already familiar from Chapter 4, which could be said to imply that death is worse the more expected life is deprived the person. The ageist implication of the deprivation account is that any life that a person is deprived of by death has positive value; therefore, it is more wrong to kill a younger person than it is an older person, other things being equal. Conversely, under the deprivation account, it is less wrong to kill a person with less prospect of continued life than it is to kill a person with a greater prospect of continued life, other things being equal.[58] I will attempt to show that both these views, which can be used to justify denying older persons life extension, are unjustly discriminatory towards older persons and therefore unsuitable as a guide to fairly distributing life extending technology.

8.1.1 Is the fair innings claim really a claim about justice?

We immediately notice that the fair innings argument, at least the way it has been articulated in the literature, is cast largely in the language of justice. Every person is "entitled" to reach a fair innings; by not reaching a fair innings an "injustice" is done to the person.[59] But one might doubt whether it is a matter of injustice that a person does not reach a fair innings. Michael Rivlin (2000) has suggested that talk of having had a 'fair share' of life might simply be understood in terms of it being 'fortunate' or 'lucky' that a person lived a certain number of years, or 'unfortunate' or 'unlucky' that a person did not live a certain number of years, without necessarily deferring the issue to the pressures of egalitarian justice. This interpretation might be viewed as a libertarian response to the welfare egalitarian implications of the fair innings argument. It is

[58] For discussion of a similar idea see Lippert-Rasmussen 2007.

[59] Notice from my earlier discussion of egalitarianism, that egalitarian justice demands that the "bonus years" of those living beyond the threshold of a "reasonable" lifespan are tradeable in order to help younger people reach the threshold.

libertarian in the sense that it implies that nature, in regard to lifespan, is inherently just and needs no help from us to level the playing field. If you live longer than someone else that is simply your 'good luck' and there is no requirement of us to redistribute resources to those who have not been lucky enough to live as long as others. Interpreting the fair innings claim in this way provides a way to look at the fair innings debate through the helpful framework of the major theories of distributive justice. I will likewise recast the fair innings argument itself in both utilitarian and egalitarian versions in order to see if this helps us identify the fundamental point being made by the argument.

Alternatively, the response could be interpreted as saying that the older person may have 'earned' their long life, for example, by exercising a healthy lifestyle, whereas the younger person may not have taken care of themselves or treated their life with the same respect, thus it is not always obvious that we should favour the younger person over the older person when it comes to distribution of medical resources, and, in fact, it may be seen as unjustifiable to redistribute resources in such a way as to bring those who have not had a reasonable lifespan up to a fair innings if they have not earned a prolonged life.

The libertarian aspect of the interpretation of fair innings as luck or chance can also be drawn from Rivlin's claim in the quotation above about something being unfair or unjust only when a human agent brings it about. In regard to the life already lived by the older person, it seems as if the fair innings argument is saying that life and the benefits therein were somehow given to the older person by some agent in the sense that the older person was given, at some time during their life, a larger piece of cake than the younger person, so to favour them again with yet more cake would be unfair to the younger person who had not yet received their first share of the cake.

But this way of talking does not always make sense in regard to 'life' or 'wellbeing'. Think about when the older person was younger, assuming he did not receive any life-saving intervention in the past, would we want to claim that at that previous time the person (the one who is now older) had not yet been given something that they now have been given by some agent? Or would we simply want to say that he had traversed his life successfully in order to get all the way to seventy years of age? It may be right then to suggest that before we can legitimately debate a welfarist redistribution of a 'fair share of life' it must be shown why the life one has is not simply a result of chance or what one has earned, the vicissitudes of a free market, as it were.

The libertarian response to the 'fairness' implications of the fair innings argument is that it would be unjust to give an artificial helping hand to the person who has not yet reached a 'fair share of life' simply because there is such a thing as a fair share of life, in order that this person may get the chance to enjoy something that an older person got to experience largely through chance or through their own effort. But I will set aside the libertarian interpretation of fair innings, as it is not relevant to the central question of this chapter.

8.2 Different moral interpretations of fair innings

I have suggested above that the fair innings argument itself might be recast into utilitarian and egalitarian versions in order to help us get at the essential point of the argument that seems so intuitively compelling.

A utilitarian interpretation of the fair innings argument is that, after a 'fair innings' has been reached, older persons lose all right to life because after a certain age, people are too costly to keep alive. The benefit they experience from continued life is not worth the opportunity cost of the resources used in keeping them alive. The utilitarian interpretation of the fair innings argument, in order to fully commit to the view, would have to include some kind of 'compulsory euthanasia' at the end of a fair innings, as this would result in the best economic outcome overall by relieving the burden on society of those members continuing to live beyond a fair innings.[60]

In contrast, the egalitarian interpretation of fair innings is that, after a 'fair innings' has been reached it is justifiable, in the context of saving lives, to redistribute medical resources in favour of younger persons in order to ensure they reach a fair innings as well. This does not imply that the continued life of an older person who has reached a fair innings has no value at all on the egalitarian interpretation, just that, in situations where we need to make a trade-off between the two lives, we should give priority for life prolonging resources to the younger person.

The utilitarian and egalitarian versions of the fair innings argument can also be viewed in terms of how each treats persons. The utilitarian version of fair innings treats

[60] cf: "Logan's Run" scenario. The utilitarian version of the fair innings argument may also be appreciated as another version of the 'overpopulation objection' to prolonging lives, i.e. there are a fixed number of years we are entitled to live beyond that we give up our right to life.

the life of a person 'instrumentally'. A person is a passive thing and the value of a person's life lies in how far the available life might be utilised to produce the outcome with the greatest overall wellbeing. Thus a fundamental right to life lasts only up to the threshold of a 'fair innings' and then evaporates when that life no longer adds to the net utility in the world. The egalitarian version of fair innings treats the life of a person in a slightly more 'fundamental' way, but still as a means to satisfying an abstract moral principle. This moral principle is to make people equal in wellbeing and this means equal in how much life one is allowed to experience. In its effort to raise the level of wellbeing of existent person to parity, its goal is to redistribute wellbeing or continued life from those with a surplus of it to those deficient in it. But in doing so it is promoting wellbeing and continued life for a certain group of people, while not strictly doing so for another group of people.

One might have thought that under a Rawlsian idea of protecting fair equality of opportunity for life extending interventions we must give priority to the young over the old because the old have already had a 'fair innings'. However, in a choice between providing a young person with the life extension pill and providing an old person with the life extension pill, we should give priority to the person least advantaged with respect to life expectancy, and this might sometimes be the older person. Fairness in the distribution of continued life might sometimes mean treating the old differently from the young and giving them priority because their time is running out for more life. The young person will, other things being equal, be an old person in the future and thus might be given priority for life extending interventions when that day comes.

This idea of treating people relative to the stage they are at in the lifespan is the fundamental idea behind Norman Daniels's 'prudential lifespan account' of healthcare distribution. The prudential lifespan account says that justice between age groups is not a problem of competition for resources between groups at a moment in time, but rather resources should be allocated to protect opportunity for healthcare at each stage of life. Therefore, just healthcare distribution ought to take an entire lifespan view. The basic idea, says Daniels, is that "since we all age, we should take as a model for what is fair between age groups what it is prudent for us to do for ourselves at each stage of life" (2008: 162). Ageing is something we all do and thus is a factor different from other considerations such as race or sex for the purposes of distributive justice. We do not necessarily create inequalities if we treat the young and the old differently, "if we treat the young one way as a matter of policy and the old another way, and if we do so over

their whole lives, then we treat all persons the same way" (2008: 171). What Daniels seems to be saying here is that, the old were once young, and the young will be old some time in their lives, so if we treat 'young people' always a certain way when it comes to policy matters, and 'old people' always a certain way, then this is treating all people the same throughout their entire lifespan. The young and old persons should be viewed as the same persons but at different stages of their lives.

Daniels's prudential lifespan account could be seem as supporting my interpretation of the Rawlsian difference principle, especially when it comes to distributing life extending interventions.

As a result, inequalities of opportunity resulting from age are not necessarily unjust; in fact, they may have benefits for everyone, because disparate treatment of people by virtue of age, when under a whole lifespan view, is really treating them equally over their entire lives. So the idea behind what Daniels calls his Prudential Lifespan Account of justice between age groups in the distribution of resources is to view the young and the old not as two groups in competition between each other for scarce resources, but rather two groups representing different stages of our lives, thus the problem of distribution is allocation of resources throughout the duration of a life. Paradoxically, only by treating people differently is 'equal access' to a resource being fulfilled.

The fair innings proponent might argue that the fair innings argument does, in fact, give priority to those with the lowest life expectancy, because, in the relevant cases where fair innings is employed, there are usually two people, one young and one old, *both* who without intervention will soon die, both in competition for scarce medical resources, so whose life should we prolong? In such cases where projected life expectancy without intervention is equally low, the fair innings argument says that intervention should go to the person who has benefited the least in terms of life years already lived, which will always be the young person, and that this is what is fair.

Daniels argues that the prudential lifespan account of justice between age groups is the account of justice people would rationally choose to adopt because they would be giving equal consideration to their needs at all points in their life. Daniels suggests that, were we to pretend not to know how old we are, that we are under, say, a Rawlsian 'veil of ignorance' and did not know what age group we would occupy in society, it would

be rational for us to make sure that we had equal opportunity for access to healthcare at all stages of our life, relative to the needs and risks of a young or an old person.[61]

Colin Farrelly (2010) has suggested that Daniels's Prudential Lifespan Account of just healthcare distribution in fact shows why life extending interventions, such as retarding ageing, ought to be given more priority than they do in the current distribution of funds and healthcare resources, namely, because the goal of retarding ageing has health benefits throughout the adult lifespan, across many stages of our lives. If we want to ensure that people benefit equally in regard to equality of opportunity to healthcare, a more efficacious way of bringing this about is to direct our goals for medicine towards curing ageing, and along with it age-related risk of disease and debility, rather than directing funding and resources towards trying to address health inequalities by curing one disease at a time. He suggests that, "rather than identify some of these forms of disadvantage as ones justice requires us to redress (i.e., those that arise due to disease) and others (i.e., age-related disadvantage) as beyond the pale of justice, we should focus our energies on what really matters to every living person, young and old alike—*removing the obstacles…to extended healthy living*."[62]

8.3 Should a 'fair innings' be determined by years lived?

There are two further presuppositions in the fair innings argument that need to be examined. The argument presupposes that (a) the value of a life is derived from how many years a person has already lived, thus when a person reaches seventy he has had his chance to experience a "fair share" of life;[63] and (b) that the value of a life can also be derived from a person's proximity to death, thus any "bonus" years beyond seventy are less valuable because there will inevitably be less and less of them in the future.[64]

An opponent of fair innings can respond to presupposition (a), which says that the number of years already lived is a good indicator of who has had a 'fair share' of

[61] See Battin 1987 for similar view of healthcare distribution.

[62] Farrelly 2007: 516. Also, see Overall 2003: 205-217 discussion of distribution of life extension between age groups.

[63] This is the same as interpreting the fair innings argument as saying that the right to life only covers life up to the quota but the entitlement disappears the moment the person lives any longer.

[64] Again, this can be expressed in right to life terms as, a person's right to life diminishes the closer her perceived proximity to death, i.e. any life beyond the fair innings threshold means less possibility of continued conscious experience, translating into that life being less valuable compared to the life of a younger person below the threshold.

life, by suggesting that is not necessarily true that just because one is chronologically older one has had more opportunity to experience the goods of life than one who is chronologically younger. Certain contingent circumstances inherent to a particular life, e.g. economic circumstances, chronic medical conditions, personality, etc., may inhibit a person from accumulating valuable experiences, even into old age. Individual lives differ greatly in their experience-to-life-years ratio. It only makes it likely that just because one has lived many years that one has experienced many of the goods of life, but this is not necessarily true. Daniela Cutas (2008: 2) has highlighted this problem in counting life years lived:

> Deciding on how much life is enough is likely to involve more injustice for people who, although they have lived for many years, have had fewer opportunities and fewer chances to flourish. The value of persons' lives cannot be dependent upon the quantity of life at one end or the other, but on their respective capacities to value their own lives, whatever the content of that which they actually value.

The defender of fair innings, however, might simply concede this point about qualifying what exactly they are talking about when they claim the older person has accumulated more of than the younger person, namely, it is a quota of 'life-experience' rather than simply quota of years. But can the fair innings proponent really make this move? This brings up the point about distinguishing between the value one gives to one's own life, and the value of that persons' life, or their life-experiences, objectively speaking. When evaluating a life with an eye for who has 'lived' the most life, it is rarely the case that 'all other things being equal' is true from the person's own point of view. The depressed person might argue she has hardly been allowed to live despite her accumulating age, while the happy person might concede he is content with the brief number of years he has been blessed to experience.

Consequently, I would argue that the claim regarding the real-life discrepancy in the experience to life-years ratio between different people at different ages still carries some weight against the idea of a fair innings. This is because the fair innings argument will always emphasise the value of length of life rather than the quality of life. The argument aims to limit the length of life if this is what intergenerational equity demands but it does not take account of quality of life lived.

There are a couple of responses we can give to presupposition (b) which says that the value of one's continued life diminishes the closer one's perceived proximity to death. The thinking here is that the further beyond the threshold of a reasonable

lifespan, seventy years, the less valuable any extra years lived. This can also be interpreted as saying that the closer one is to death not only is one's desire to live lessened but the less valuable that extra bit of life left is to the older person and to society. One might start by questioning the relevance of life expectancy as a criterion for survival. It could be argued that so long as the older person wishes to go on living, whether this ends up being one year or forty years, their continued life is as valuable as the life of a younger person. Regardless of how short a life expectancy, this life can be valuable to the person who desires life equally as strongly as the person with a greater life expectancy. This response relies on the value of the continued life of the person to the person himself or herself.

But a more objective response can be given to the claim about proximity to death. The possibility of indefinitely prolonging one's lifespan, in a sense, discounts the ageist proximity to death presupposition because, presuming it does not matter when one starts the life-extension treatment and that this treatment can go on indefinitely, the expectation of continued life is brought to parity between the chronologically young person and the chronologically old person. Given life extension, it would no longer be true that older people must consider living beyond 70 years as "a sort of bonus beyond that which could reasonably be hoped for" (Harris 1985: 91).

Farrant (2009) has argued that one implication of the existence of life extending technology for the fair innings argument is that, if this technology did exist, then the fair innings argument would be limiting the lifespan of individuals. The implication is that what the fair innings argument is doing is establishing a normal lifespan rather than work without regard to the empirical realities. The establishing of a reasonable lifespan at seventy is an out-dated notion of an acceptable lifespan.

The fair innings proponent may continue to object on the grounds of justice even in the radical longevity situation. He may argue that if life extending technology were in limited supply, and we were faced with a clear situation of prolonging the life of 20 year old or an 80 year old, we should still favour the younger person over the older person in regard to access because of the first claim we identified, viz., that the value of a life is derived from how many years a person has already lived, therefore it would be unjust to prolong the life of the 80 year old over the 20 year old. Furthermore, given the possibility of radical life extension, a greater injustice is done the younger person if they are denied access to it over the older person.

But it seems that in one sense the same degree of injustice would be done both to the younger and older person in this longevity-enhanced world, because both would have had, in principle, the same potential for continued life. So there is no reason to favour the younger person over the older one. However, the defender of fair innings could still say that an older person who has had a fair innings, perhaps now five-hundred years old, has less right to continued life than the younger person simply because he had already had more life. Nevertheless, in the world of radical life extension, the very notion of a falling short of a 'fair innings' is made much less meaningful: it is less compelling to argue that an injustice has been done to the four-hundred-year-old who is denied access to life extending technology over the five-hundred-year-old. Of course, this is all based on speculation about how such a longevity-enhanced society would work and the nature of the longevity technology.

One final point can be made against applying the fair innings argument to the distribution of radical life extending technology. The situation in which radical life-prolonging medical interventions, of the kind we have been discussing in the thesis, are given to the older person in preference over the younger person, would rarely arise because, by their very nature, these life extending interventions are the kind of medical intervention needed only by older people to postpone death from old age and age-related damage, and thus are not needed and would not be used by younger people. But this is to a large extent an empirical question.

It might simply be conceded that the stronger claim of the fair innings argument, i.e. that giving the older person a life-prolonging intervention in favour of giving the younger person the same intervention would be unjust, is compatible with the view of the conditional nature of the positive right to life I described in earlier sections, i.e. that a person's positive right to life might clash with the similar right of another person in certain circumstances.

8.4 The deprivation account (again)

I have assumed the truth of a certain version of the deprivation account of death's badness throughout the thesis (see Chapter 4). The event of a person's death is bad for them because this event deprives them of the goods of life and also the opportunity for

more. The deprivation account of death's badness, however, might be seen as another kind of fair innings claim and thus used as another justification for the argument that older persons deserve continued life less than younger persons. The deprivation account could imply that a person's death is worse the more life is deprived the person at the time of their death, other things being equal. Consequently, it would be more wrong to allow a younger person to die than it would an older person by depriving them access to life extending interventions, because the younger person has a greater potential for continued life.

However, the deprivation event view doesn't imply that a person's being dead continues to be a bad thing for them after the event of death. But this being said, the view might allow that the harm of the event of death can be of lesser or greater magnitude to the person who dies, lesser harm to the person who is satisfied with the life or had the opportunity to enjoy the benefits of life that they've had up till the point of their death[65], greater harm to the person who dies before they have had the opportunity to experienced the goods of life or satisfy certain long-term desires, and so on. But no harm continues to accrue to the person after they die. So what we are talking about here is whether or not it makes sense to say that some events of death are worse than others. To say this we need to go return to the idea of the state of being dead being comparatively worse (other things equal) to state of being alive. On this view he badness of a person's death derives from comparing the state of being dead to the state of being alive, of continuing a conscious life and all the potentially goods things to be experienced therein. The latter obviously being a much better state of affairs for the person.

What does this mean for our discussion of the supposed asymmetry in our attitudes towards the continued life of younger person compared to older persons? The deprivation state view (but maybe not the deprivation event view) suggests that it would be more wrong to allow to die the person who has more conscious experience to lose than the person who has less conscious experience to lose, other things being equal. And the younger person has more life to be deprived of than the older person.

Let's compare the responses of the utilitarian and the Kantian on the issue of whether some deaths are worse than others and maybe they will shed some light on the

[65] The idea of a premature death is not necessarily guided by the number of years lived by the person at the time of death, but could also be influenced by the amount of opportunity to experience the goods of life, or the capacity to enjoy the benefits of living, etc., that the individual has been allowed.

problem. It looks like the utilitarian would have to accept the implications of the deprivation state account, that death can be bad even if it is not a harm, this is because they are committed to interpreting the loss of continued conscious experience as something that can be given an amount or quantity: the outcome in which the younger person is killed or allowed to die instead of the older person is worse because the loss to the younger person in terms of continued life is larger than would have been the case if the older person had been killed or allowed to die. The preference utilitarian, for example, might argue that the younger person's death was a worse outcome because its frustration of the person's preference to continue living was greater than the frustration of the same future-orientated preference that would have occurred had the older person died instead.

The Kantian, on the other hand, might have a different interpretation of the deprivation state account of death's badness. The Kantian might claim that a persons' death is bad not because it deprives the person of continued conscious experience, which is comparatively better than being dead, but that it deprives the world of something intrinsically valuable, namely, autonomy or a rational nature, as it is instantiated in the life of a person. In this sense, the 'death' of human agency or the rational nature inherent in the life of the person, is the badness done by killing a person. But for the Kantian, because what death deprives the person (and the world) of is the capacity for autonomy and agency, there is no difference, other things being equal, between the deaths of a young person and an old person, both still full possessors of the valuable features of rationality and autonomy. All deaths of autonomous persons are equally bad. From the Kantian perspective, allowing to die either a young or an old person is equally bad, other things being equal, because both deaths deprive the world of the existence of moral agency.

If we accept a utilitarian interpretation of the deprivation account we must see that killing or allowing a younger person to die would be worse than killing or allowing an older person to die, because the badness of the younger person's death is greater than that of the older persons death in that the former are typically (though not necessarily) deprived of more conscious experience than the latter. However, if we accept a Kantian interpretation of the deprivation account, we see that what the deprivation account actually helps to show is that it would be equally bad to allow to die the older person as it would be to allow to die the younger person, because both deaths entail the equal loss of a valuable feature of the life of a person, namely, autonomy or a rational nature.

The deprivation account differs from the fair innings argument because it only seems to give weight to the value of potential future life that has been deprived a person, but seems to give no weight to how much life the person has already lived. The wrongness of killing or allowing a person to die on the deprivation account is guided solely by how long the person would have had left to live. Thus one response to the deprivation account is similar to one of the responses I gave to the fair innings argument, that given the possibility of radical life-extension we can no longer claim that the older person has less life to lose than the younger person. This is because the possibility of successful life extending interventions opens up the possibility of indefinite extension of conscious experience, a possibility that, in principle, brings the possible future life or future conscious experience of the young person and the older person to parity. Given this possibility, it would be equally wrong to allow the older person to die as it would the younger person, as they both, in principle, have an indefinite potential for continued life.

8.5 Conclusions

In this final chapter I have argued that the fair innings claim and the Nagelian deprivation account are views unjustly discriminatory against older persons in respect to continued life and therefore unsuitable as a guide to a fair distribution life extending technology. This result coincides with my conclusions from Chapter 6, namely, that (among those whom a life extension pill will benefit) a life extension pill should be distributed equally unless giving priority to those with lowest projected life expectancy makes them better off. Older persons and younger persons may alternatively be members of the least advantaged group in respect to life expectancy, i.e. the opportunity for continued life. Therefore my conclusion is that denying life extension in principle to older persons when the opportunity for more life would benefit them is impermissible.

CONCLUSION

What has this thesis concluded? At least this: We would not be acting impermissibly by developing life extending technology and extending human lifespans beyond their current limitations. There are deep challenges to the benefit of humans living substantially longer lives but they are certainly not indomitable ones. While many objections to life extension have urgent intuitive appeal (e.g. that what it means to be human would be cheapened or lost entirely if we had longer lives than we currently have; or that having life extension would inevitably lead to disastrous overpopulation and there would be no positive alternative to the threat of overpopulation) they often miscarry on closer examination. Others present a more determined challenge to the benefit and permissibility of having longer lives and/or having life extending technologies (e.g. that lower wellbeing would inevitably result from the introduction of life extending technology into our already overpopulated and resource limited world; or that an unequal access to life extending technology would be a worse fate than never having had such technology at all), but these challenges can be met with solid responses that disprove accusations of impermissibility.

I do not go so far in this thesis as to claim that life extension should be pursued, but I do conclude that life extending technology is permissible, which the reader will remember from chapter 1, was the sole aim of the thesis. Nevertheless, my conclusions might be seen as supporting transhumanism, the view that humans should be permitted to use any technology to modify or enhance human mental or physical capabilities. I have argued that developing life extending technologies and enhancing the human lifespan is not impermissible, but this is not exactly a full-throated endorsement of transhumanism. The difference between my thesis and the transhumanist's is that there is a certain kind of transhumanist (e.g. Bostrom; de Grey) who wants to prioritise life extension above all other increases to life expectancy. They are zealously utilitarian in this, as they believe we ought to promote increases to wellbeing through extension of the lifespan without limit. My view, on the other hand, is not a utilitarian one in this respect. I have argued (see Chapter 6) that we ought to prioritise life extension to those less well off in terms of potential life expectancy, giving them as much chance at a longer life as resources allow. You don't need to be a radical utilitarian to think that life

extension is good, at least for those existing people who desire to prolong their own lives. My conclusions in the thesis imply sympathy for the benefit that might be gained from human enhancement technologies and at least a provisional approval of how these technologies may shape the future of humanity.

Far from being a sign of something negative, a bad omen for humanity's future, people living longer lives and having children less frequently are the consequence of a species undergoing an increasingly sophisticated existence, one whose ultimate teleological finale might be a yet more refined and distilled version of our species. Individuals living longer lives is a by-product of such things as better healthcare and education, advanced cures for human sufferings and increasing comforts, and a sign of less desperation to populate the Earth with new members who will have to start over from scratch. Humans with longer lifespans is an indication that, as a species, we are on a trajectory moving further and further away from the short, brutish lives that were our origin.

APPENDICES

A Basic assumptions about immortality

The notion of an infinite life is fundamentally different from that of an artificially extended but finite life. This difference is not just a difference in order of magnitude with respect to number of years lived. To be necessarily immortal is to be in a state without death, not merely a state in which one did not age or decline from the inside. It is to never have to face non-existence, to exist always, whether this was once an embodied survival or a non-embodied survival or some other as yet unfathomed kind of survival. As such, it is difficult to know what to say about such a state, especially for a conscious human life.

It is often unclear what writers mean by 'immortality' both in the philosophical literature and in general fiction and non-fiction. A never ending life is an abstract notion when it comes to the life of a human being, a condition involving much speculation and extrapolation, rather than statements of empirical realities, in attempting to describe it in any detail. However, frequently writers have simply passed over any clarification of what kind of thing they are referring to by immortality in respect to a human life. In this section I will clarify what I mean by immortality and what it means for a recognisably human life. I will outline some basic assumptions that I take for granted when talking about immortality, as well as human embodied and non-embodied existence.

'Necessary' and 'contingent' immortality

From the outset we need to distinguish between what I will call 'contingent immortality' and 'necessary immortality'.[66] Most people when they hear the term immortality imagine a kind of contingent embodied immortality, i.e. they are thinking of an individual who is perhaps ageless, and therefore has the potential for great longevity by virtue of being exempt from disease and decline associated with ageing, but who could still have their existence cut short by external bodily harm, either from

[66] The distinction between contingent body-bound immortality and necessary body-bound immortality was first introduced by Hunter Steele, 1976.

the actions of others (e.g. homicide) or themselves (e.g. suicide) or by non-agential external events (e.g. natural disaster). Thus contingent immortals are not invulnerable or exempt from death. On the other hand, such an individual would not qualify as having necessary immortality, which can literally be defined as a state without death. The necessary immortal will never have to not exist.

This everyday misunderstanding of immortality is understandable, for when we talk casually about immortality for a human life, that is, something like an embodied existence of very long length, and not the kind of non-physical immortality of a self or soul in religious contexts, it looks as if we must always mean contingent immortality, or a 'near' or 'virtual' immortality. Although we may one day cure ageing in its entirety and, as a result, no longer suffer bodily ageing or decline, we can still expect our lives to be cut short by contingent external factors. Put another way, when we cure ageing it will be true that we cannot reasonably expect to die 'from within', but external violence will still play a role in the limit of our potentially endless embodied existence. Thus it is unlikely that 'immortality', in the strict necessary sense of the term, could ever be our human condition or the condition of any other living embodied thing.

Clarifying this basic ambiguity in the term immortality tidies up a whole range of less useful argument that would have been directed at the notion of humans leading a necessary immortal life, i.e. an indestructible and infinite existence.[67] For my part, I believe that the notion of necessary immortality is less useful in any discussion of whether an immortal lifespan would be desirable or undesirable, good or bad, for a human person. Nor, I would suggest, is invulnerability or indemnity from death a state even desired by those who desire a radically longer lifespan or 'immortality'.

Embodied and non-embodied human existence

I have suggested that an aspect of the commonsense view of human immortality is that of an embodied existence. But there has also been speculation about potential forms of non-embodied immortal existence that may be available to humans in the future, usually made possible by new technologies. By 'humans' I mean available to our continuing

[67] Burley, on the other hand, disagrees about any real benefit gained from the distinction: "Since I suspect insuperable conceptual difficulties attach to each of these options, I am doubtful that deciding whether it is contingent or necessary immortality that is at stake will result in the debate's being placed on a firm footing", 2009: 541.

human 'consciousness' or 'identity'. For example there has been discussion, most noticeably in fiction but also in the academic literature, of living a never-ending existence as a consciousness uploaded into a computer in the manner similar to that depicted in Gibson's *Neuromancer*.[68] This kind of non-embodied life extension clearly highlights the issue of whether one will still be living a recognisably human life given such technological non-embodied existence.

The posthumanist it claims that non-embodied existence, the transcending of the limitations of the human embodied existence, has always been the ultimate end for human life. The goal for humanity, says the posthumanist, is to achieve immortality in the name of human advancement. And one potential way of achieving this is by the survival of the human consciousness through technology. However, as the objection goes, it is questionable whether non-embodied consciousness would indeed be humanity surviving into the future. Some might think it an obvious problem,[69] but why must having potentially no bodily or temporal limits to existence entail a loss of being human? To take one example, if we did not have the sense organs we are so used to in our embodied existence, would we still be able to send and receive stimulus from the outside world? If we were digitalised and tied to some kind of hard drive, we might be able to send and receive various data through input and output devices connected to the hard drive. But how would we think? Is that, perhaps, all we would be, a purely 'thinking' thing? One thing is clear, that this will be an existence completely different from the embodied one humans now experience.

Could there be such a thing as the death or end of one's consciousness or personality as there is of one's body? Just as there are limits on the human body are there necessarily limits on human consciousness? This is probably a contingent empirical matter. If one's conscious life could be prolonged indefinitely by technology in the future, for example, by being digitally copied over and over again and moved around among whatever available storage devices, then there seems no reason why a 'saved consciousness' could not conceivably continue on without limit, the limitations obviously being the limitations of available storage devices in existence.

It seems both an embodied and non-embodied immortal could have their existence ended. Even with forms of non-embodied existence, immortality would not be

[68] See: Astakov 2007; Clark 1995: 39-52, for examples of academic discussion of potential non-embodied existence.

[69] Some commentators have made it clear that any change in what we think of now as 'human nature' would be just as good as extinction; see Callahan 1995; Kass 2004.

certain, but would always be contingent. Whatever technological hardware one's consciousness was tied to would eventually degrade over time. Admittedly, though, this would be a very long time, and we must assume there would always be the opportunity to 'copy' the 'data' that is one's consciousness onto a fresh hard drive or whatever the technology may be.

It is an open question whether (1) non-embodied existence will continue to be recognisably human existence, and (2) if it is an existence that is no longer human in the most basic ways, then would this be just as bad as human extinction, or simply another (albeit radical) stage in human evolution?

Partial and complete perspectives on an immortal life

It is sometimes said that, on a particular sense of self, one can take a view of one's life 'as a whole'. Is it possible to have the same perspective on one's potentially never-ending existence as it is for a necessarily finite existence? If not, how can we gauge whether an infinite existence is desirable or not? Mikel Burley explains the argument, "if one agrees that a necessary condition of being able to assess the desirability of a life is that the life can be conceivable as a whole, then it looks as though such as assessment cannot be made in the case of a putative immortal life" (2009: 539). But is it really a necessary condition for the desirability of a life to be able to conceive of it in its entirety? Take a different example. There are plenty of open-ended activities, i.e. that we pursue and that we can conceivably never see ourselves coming to the end of, but this does not make a life that was made up with such activities undesirable just because we cannot conceive of it in its entirety. I would suggest that there is similarly no requirement that an immortal life be viewable as a whole that should influence its desirability.

Stephen Clark makes a comments that puts into perspective the idea of conceiving of an immortal life as a whole: "Perhaps something that is practically indistinguishable from pure immortality is possible after all: living for an extra ten millennia is indistinguishable now from 'living forever'".[70] Clark is right, from the perspective of the individual starting out on a contingently immortal life, ten thousand

[70] Clark 1995: 11-12.

more years is just as good as 'necessary' immortality. From the contingent immortal's present perspective, it will *always* be just as good as having necessary immortality, because, as far as they know, they may never reach an end of it.

On an even more abstract note, if our existence were one of a non-embodied consciousness, would we be able to accommodate into the narrative of our lives an ending, or any structure at all? As a non-embodied personality, would our identities be more threatened over such a great length of time, or would it not be necessary for us to seek some kind of limiting narrative for our non-embodied lives? I think the same thing would apply to contingent non-embodied existence as it does to contingent embodied existence, namely that, there would be no requirement that such a life be conceivable as a whole in order for it to be thought of as desirable or undesirable.

A healthy and active immortality

One important assumption that must be made when discussing the desirability of immortality is that individuals who opt for the 'immortality intervention', if it were ever to become available, would be at least satisfactorily healthy and energetic and would continue in this condition across the span of their indefinitely long life. As Clark again comments:

> On the one hand, living forever is plainly intended to be living healthily, or youthfully, forever (with whatever caveat about youth and health proves necessary)…It would be too easy to make immortality sound worthless if we guaranteed unending pain, or deprived the immortal of her wits. The question is: would living forever be essentially a bad idea, even if all ordinary advantages of health and good sense were added?[71]

So the same qualification, which must be assumed about moderate life extension, should also be assumed when discussing immortality as well: that it will not be a life of less than average health and vigour, but a life where the individual is of at least average if not better wellbeing and healthiness throughout the majority of it.

[71] Clark 1995: 12.

Immortality and isolation

A criticism that has been levelled at those drawing conclusions about the desirability of immortality from fictional, e.g. Williams's use of the Makropulos play, in which a lone female immortal lives a life of loneliness and boredom, is that necessary isolation as a necessary part of immortality is highly in the event that the 'immortality intervention' ever becomes available. More likely is that there will be many immortals living in the world sharing the same situation, maybe even one's own friends and family. So many of the problems directed at living an immortal life that relate to boredom and desire are influenced by fictional cases where the immortal is left in isolation. Thus, similar to the assumption above about the immortal expecting to live in relative good health and vitality, so too can they expect to live in a world peopled by other immortals, with whom they can share their unprecedented existence.

Immortality and limited resources

This is a problem for societies and the world at large given a population of immortal individuals. By all accounts it is an intractable problem, for any society must necessarily be limited in resources, as must our current planet. However, it could be speculated that eventually the immortals existing in any society will live long enough to see the possible colonisation of other planets for their resources, and therefore they may outlive the resource problem here on earth. This kind of speculation about colonisation of other planets may be the only way of avoiding what looks like a stubborn problem for the desirability of making many people immortal who will live in a world of limited space and resources.

There is also another possible way out. The prospect of non-embodied immortality—one's life as a 'consciousness' tied to some piece of technology—might allay worries about the planet's limited resources. The prospect of mind uploading, for example, of abandoning the flesh and all the resources it must consume to survive, would definitely conserve space and resources for a global population of potential immortals. Of course, some space and power resources will inevitably be needed to

sustain this non-embodied existence for all those consciousnesses tied to hard drives, but not nearly as much as for a global population of embodied human beings.

B List of figures and tables

BIBLIOGRAPHY

Ackerman, Felicia N. 2009. Death is a punch in the jaw: life-extension and its discontents. In Bonnie Steinbock. ed. *The Oxford Handbook Of Bioethics*. Oxford: Oxford University Press, pp. 324-348.

Aristotle. *Nicomachean Ethics*. Cambridge: Cambridge University Press, 2000.

Armstrong, R.J.E. and Svendsen, C.N. 2000. Neural stem cells: From cell biology to cell replacement, *Cell Transplantation*, 9: 139–152.

Astakov, V. 2007. The continuum of consciousness: mind-uploading and resurrection of human consciousness. Cornell University Library, <http://arxiv.org/ftp/arxiv/papers/0710/0710.2556.pdf>.

Battin, M.P. 1987. Age rationing and the just distribution of health care: is there a duty to die? *Ethics*, 97: 317-340.

Benatar, D. 2006. *Better Never To Have Been: The Harm of Coming Into Existence*. Clarendon Press: Oxford.

Bentham, Jeremy. [1789] 2000. *An Introduction to the Principles of Morals and Legislation*. Batoche Books.

Blackford, R. 2009. Moral pluralism versus the total view: why Singer is wrong about radical life extension, *Journal of Medical Ethics*, 35: 747-752.

Bodnar, A.G., Ouellette, M., Frolkis, M., Holt, S.E., Chiu, C., Morin, G.B., Harley, C.B., Shay, J.W., Lichsteiner, S., Wright, W.E. 1998. Extension of life-span by introduction of telomerase into normal human cells, *Science* (New Series), 279: 349-52.

Bortolotti, Lisa and Nagasawa, Yujin. 2009. Immortality without boredom, *Ratio* (new series), XXII 3: 261-277.

Bostrom, N. 2005a. The fable of the Dragon Tyrant, *Journal of Medical Ethics*, 31: 273-277.

Bostrom, N. 2005b. Recent developments in the ethics, science, and politics of life-extension, (forthcoming, *Aging Horizons*, Sept/Oct issue, 2005).

Bostrom, N. 2009. The future of humanity, *Geopolitics, History and International Relations*, 1: 41-78.

Broome, J. 2005. Should we value population? *Journal of Political Philosophy*, 13: 399-413.

Bunn, A. 2009. Evaluating life extension from a narrative perspective, *American Journal of Bioethics*, 9: 79-80.

Burley, Mikel. 2009. Immortality and meaning: reflections on the Makropulos debate, *Philosophy*, 84: 529-547.

Callahan, D. 1995. Ageing and the life cycle: a moral norm? In Callahan, D., Meulen, R.H.J., Topinkova, E. eds. *A World Growing Old: The Coming Health Care Challenges*. Washington, D.C.: Georgetown University Press.

Caplan, A. 1981. The unnaturalness of aging: a sickness unto death? *Concepts of health and disease: Interdisciplinary perspectives*, 725-737.

Caplan, A. 2005. Death as an unnatural process, *EMBO reports*, 6: S72-S75.

Carnes, B.A., Olshansky, S.J., Grahn, D. 2003. Biological limits to the duration of life, *Biogerontology*, 4: 31-45.

Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Chen, K.S., Masliah, E., Mallory, M., Gage, F.H. 1995. Synaptic loss in cognitively impaired aged rats is ameliorated by chronic human nerve growth factor infusion, *Neuroscience*, 68: 19-27.

Clark, Steven. R.L. 1995. *How To Live Forever: Science Fiction and Philosophy*. London: Routledge.

Cutas, D.E. 2008. Life extension, overpopulation and the right to life: against lethal ethics, *Journal of Medical Ethics*, 34: e7.

Daniels, N. 2008. *Just Health: Meeting Health Needs Fairly*. Cambridge University Press.

Davis, J.K. 2004. Collective suttee: is it unjust to develop life extension if it will not be possible to provide it to everyone? *Annals of the New York Academy of Sciences*, 1019: 535-541.

de Grey, A.D.N.J., Ames, B.N., Anderson, J.K., Bartke, A., Campisi, J., Heward, C.B., Stock, G. 2002a. Time to talk SENS: critiquing the immutability of human aging, *Annals of the New York Academy of Sciences*, 959: 452-62.

de Grey, A.D.N.J., Baynes, J.W., Berd, D., Heward, C.B., Pawelec, G., Stock, G. 2002b. Is human aging still mysterious enough to be left only to scientists? *BioEssays*, 24: 667-76.

de Grey, A.D.N.J. 2003. The foreseeability of real anti-aging medicine: focusing the debate, *Experimental Gerontology*, 38: 927-34.

de Grey, A.D.N.J. 2004. We will be able to live to 1,000. *BBC News*. Retrieved 07/04/2008 from <http://news.bbc.co.uk/2/hi/uk_news/4003063.stm>.

de Grey, A.D.N.J. 2004b. Aging, childlessness, or overpopulation: the future's right to choose, *Rejuvenation Research*, 7: 237-38.

de Grey, A.D.N.J. 2005. Life extension, human rights, and the rational refinement of repugnance, *Journal of Medical Ethics* 31: 659-663.

de Grey, A.D.N.J. 2005. Hang in there: the 25-year wait for immortality, interview by Ker Than, *LiveScience* (online), April 11, 2005.

de Grey, A.D.N.J. 2007. Lifespan extension research and public debate: societal considerations, *Studies in Ethics, Law, and Technology*, The Berkeley Electronic Press.

Dick, P.K. [1969] 1991. *Ubik*. Vintage Books: New York.

Engelhardt , T.H. 1979. Is aging a disease? In Robert M. Veatch. ed. *Lifespan: Values and Life Extending Technologies*. Harper & Row: San Francisco: 184-194.

Ettinger, R.C.W. 1965. *The Prospect of Immortality*. London: Sidwick & Jackson.

Farrant, A. 2009. The fair innings argument and increasing lifespans, *Journal of Medical Ethics*, 35: 53-56.

Farrelly, C. 2007. Sufficiency, justice, and the pursuit of health extension, *Rejuvenation Research*, 10: 513-519.

Farrelly, C. 2010. Equality and the duty to retard human ageing, *Bioethics*, 24: 384–394.

Fischer, J.M. [1994] 2009. Why immortality is not so bad, in *Our Stories: Essays on life, death, and free will*. Oxford University Press, pp. 79-92.

Flurkey, K., Papaconstantinou, J., Miller, R.A., Harrison, D.E. 2001. Lifespan extension and delayed immune and collagen aging in mutant mice with defects in growth hormone production, *Proceedings of the National Academy of Sciences*, 12: 6736-41.

Gems, D. 2003. Is more life always better? The new biology of aging and the meaning of life, *Hastings Center Report*, 33: 31-39.

Gibson, W. 1984. *Neuromancer*. Voyager/HarperCollins: London.

Glannon, W. 2002. Extending the human lifespan, *Journal of Medicine and Philosophy*, 27: 339-354.

Greider, C.W. 1998. Telomeres and senescence: the history, the experiment, the future, *Current Biology*, 8: R178-R181.

Greider, C.W. and Blackburn, E.H. 1985. Identification of a specific telomere terminal transferase activity in Tetrahymena extracts, *Cell*, 43: 405-13.

Gruman, Gerald. J. 1966. A history of the ideas about the prolongation of life. *Transactions of the American Philosophical Society*, 56:1-102.

Harris, J. 1985. *The Value of Life*. Routledge & Kegan Paul: London.

Harris, J. 2002. Intimation of immortality: The ethics and justice of life extending therapy, *The Harold Hatch International Lecture on Longevity and Population Aging*. The International Longevity Center-USA.

Hayflick, L. 1998. How and why we age, *Experimental Gerontology*, 33: 639-53.

Hayflick, L. 2000. The future of aging, *Nature*, 408: 267-69.

Heilbronn, L.K. and Ravussin, E. 2003. Calorie restriction and aging: review of the literature and implications for studies in humans, *American Journal of Clinical Nutrition*, 78: 361-9.

Horrobin, S. 2005. The ethics of ageing intervention and life-extension. In Ratten, S. ed. *Ageing Interventions and Therapies*.

Hurka, T.M. Average utilitarianisms, *Analysis*, 42: 65-69.

Juengst, E. T. 2004. Anti-aging research and the limits of medicine. In Post, S. G. and Binstock, R.H. eds. *The Fountain of Youth: Cultural, Scientific and Ethical Perspectives on a Biomedical Goal*. Oxford University Press: New York, pp. 321-339.

Kant, I. [2012] *Groundwork of the Metaphysics of Morals*. (Trans. Mary J. Gregor). Cambridge University Press.

Kass, L. 2004. L'Chaim and its limits: why not immortality? In Post, S. G. and Binstock, R.H. eds. *The Fountain of Youth: Cultural, Scientific and Ethical Perspectives on a Biomedical Goal*. Oxford University Press: New York, pp. 304-320. Originally published in *First Things*, 113, 2001: 17-24.

Keller, G. and Snodgrass, H.R. 1999. Human embryonic stem cells: the future is now, *Nature Medicine*, 5: 151-52.

Kirkwood, T.B.L. 2002. The evolution of aging, *Mechanisms of Aging and Development*, 123: 737-745.

Kirkwood, T.B.L. and Rose, M.R. 1991. Evolution of senescence: late survival sacrificed for reproduction, *Philosophical Transactions: Biological Sciences*, 332: 15-24.

Kurzweil, R. 1999. *The Age of Spiritual Machines*. Penguin Books.

Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology*. Penguin Books.

Kurzweil, R. 2005b. Chasing immortality: the technology of eternal life, interview by Craig Hamilton, *What Is Enlightenment?* 30, 58-69.

Kymlicka, Will. 2002. *Contemporary Political Philosophy: An Introduction*. 2nd edn. Oxford University Press.

Lane, M.A., Mattison, J., Ingram, D.K., Roth, G.S. 2002. Caloric restriction and aging in primates: relevance to humans and possible CR mimetics, *Microscopy Research and Technique*, 59: 335-38.

Lane, M.A., Ingram, D.K., Roth, G.S. 2002. The serious search for an anti-aging pill, *Scientific American*, 287: 24-29.

Latham, R., trans. 1994. *Lucretius, On the Nature of the Universe*. London, UK: Penguin UK.

Ligotti, T. 2010. *The Conspiracy Against the Human Race*. New York: Hippocampus Press.

Lippert-Rasmussen, K. 2007. Why killing some people is more seriously wrong than killing others, *Ethics*, 117: 716-738.

Locke, J. 1694. Of identity and diversity, reprinted in John Perry *Personal Identity*. 2nd edn. Berkeley: University of California Press, 2008: 33-52.

Lockwood M. 1988. Quality of life and resource allocation, *Royal Institute of Philosophy Lecture Series*, 23: 33-55.

Martin, G.M. 1999. Genes that modulate longevity and senescence. In J.M. Robine, B. Forette, C. Franceshi, M. Allard. eds. *The Paradoxes of Longevity*, New York: Springer.

Masoro, E.J. 1996. The biological mechanism of aging: is it still an enigma? *Age*, 19: 141-45.

McCay C.M., Crowel M.F., Maynard L.A. 1935. The effect of retarded growth upon the length of the lifespan and upon the ultimate body size, *Journal of Nutrition*,10:63–79.

Mckie, J., Kuhse, H., Richardson, J., Singer, P. 1996. Allocating healthcare by QALYs: the relevance of age, *Cambridge Quarterly of Healthcare Ethics*, 5: 534-545.

McMahan, J. 1981. Problems of population theory, *Ethics*, 92: 96-127.

Metz, T. 2003. Utilitarianism and the meaning of life, *Utilitas*, 15: 50-70.

Mill, J.S. [1863] 1957. *Utilitarianism*. 2nd rev. ed. New York: Liberal Arts Press.

Miller, R.A. 2002. Extending life: scientific prospects and political obstacles, *The Milbank Quarterly,* 80: 155-74.

More, M. Superlongevity without overpopulation. In Immortality Institute, ed. *The Scientific Conquest of Death*. Libros en red. 2004: 169-185.

Nagel, T. 1979. Death, in *Mortal Questions*. Cambridge University Press, pp. 1-10.

Narveson, J. 1967. Utilitarianism and new generations, *Mind*, 76: 62-72.

Nozick, R. 1974. *Anarchy, State, and Utopia*. Oxford University Press.

Nozick, R. 1981. *Philosophical Explanations*. Harvard University Press.

Nussbaum, M. 1994. Mortal immortals: Lucretius on death and the voice of nature, in *The Therapy of Desire*. Princeton University Press, Princeton: New Jersey, pp. 192-238.

Oeppen, J. and Vaupel, W.P. 2002. Broken limits to life expectancy, *Science*, 296: 1029-1031.

Olshansky, S.J., Hayflick, L., Carnes, B.A. 2002. Position statement on human aging, *Science of Aging Knowledge Environment*, 24. Retrived November 20, 2007, from <http://sageke.sciencemag.org/cgi/content/full/sageke;2002/24/pe9>.

Overall, C. 2003. *Aging, Death, and Human Longevity: A Philosophical Inquiry*. Berkeley: University of California Press.

Parfit, D. 1984. *Reasons and Persons*. Oxford University Press: Oxford.

Parfit, D. 1997. Equality and priority, *Ratio* (new series),10: 202-221.

Perrett, R.W. 1986. Regarding immortality, *Religious Studies*, 22: 219-233.

Perry, J. ed. *Personal Identity*. University of California Press: Berkeley.

Pettit, P. 1993. Consequentialism. In Peter Singer. ed. *A Companion To Ethics*. Oxford: Blackwell Reference.

Pijnenburg, M.A.M. and Leget, C. 2007. Who wants to live forever? Three arguments against extending the human lifespan, *Journal of Medical Ethics*, 33: 585-587.

President's Council on Bioethics (PCBE). 2003. Ageless bodies, in *Beyond Therapy: Biotechnology and the Pursuit of Happiness*. Harper Collins Publishers: New York, pp. 159-204.

Post, S. G. and Binstock, R.H. eds. 2004. *The Fountain of Youth: Cultural, Scientific and Ethical Perspectives on a Biomedical Goal*. Oxford University Press: New York.

Rawls, J. 1971. *A Theory of Justice*. Oxford University Press: London.

Richel, T. 2003. Will human life expectancy quadruple in the next hundred years? Sixty gerontologists say public debate on life extension is necessary, *Journal of Anti-Aging Medicine*, 6: 309-314.

Rivlin, M.R. 2000. Why the fair innings argument is not persuasive, *BMC Medical Ethics*, 1: 1.

Rooskby, E. and Weckert, J. 2004. Digital divides: their social and ethical implications. In Brennan, L. and Johnson, V. eds. *Social, Ethical and Policy Implications of Information Technology*, pp. 29-47. Hershy, PA: Information Science Publishing, pp. 29-47.

Roth, G.S., Ingram, D.K., Lane, M.A. 2001. Caloric restriction in primates and relevance to humans, *Annals of the New York Academy of Sciences*, 928: 305-15.

Saunders, J., ed. 1966. *Greek and Roman Philosophy After Aristotle*. New York, NY: Free Press.

Schechtman, M. 1996. *The Constitution of Selves*. Cornell University Press: New York.

Shermer, M. 2001. Nano nonsense and cryonics, *Scientific American*, p. 29.

Singer, P. 1993. *Practical Ethics*. 2nd edn. Cambridge University Press.

Singer, P. 1991. Research into aging: should it be guided by the interests of present individuals, future individuals, or the species? In Ludwig, F.C. ed. *Lifespan Extension: Consequences and Open Questions*, New York: Springer, 1991, pp. 132-145.

Singer, P. 2010. Should this be the last generation? *The New York Times: Opinionator*, June 6.

Sorenson, R. 2005. The cheated god: Death and personal time, *Analysis*, 65: 119-125.

Steele, H. 1976. Could body-bound immortality be liveable? *Mind*, 85: 424-427.

Stein, P.K., Soare, A., Meyer, T.E., Cangemi, R., Holloszy, J.O., Fontana, L. 2012. Caloric restriction may reverse age‐related autonomic decline in humans, *Aging cell*, 11: 644-50.

Sterling, B. 1996. *Holy Fire*. Orion Books: London.

Strawson, G. 2004. Against narrativity, *Ratio* (new series), 17: 428-452.

Temkin, L.S. 2008. Is living longer living better? *Journal of Applied Philosophy*, 25: 193-210.

Trotter, G. 2004. Why bioethics is ill equipped to contribute to the debate about prolonging lifespans, *HEC Forum*, 16: 197-213.

University of Southern California. 2005. Deleting 'anti-aging' gene from yeast greatly lengthens lifespan, *ScienceDaily*. Retrieved December 1, 2007, from <http://www.sciencedaily.com- /releases/2005/11/051117181327.htm>.

Veatch, R.M. 1979. ed. *Lifespan: Values and Life Extending Technologies*. San Francisco: Harper Row.

Velleman, D.J. 1993. Wellbeing and time. In J.M. Fischer. ed. *The Metaphysics of Death*, Stanford University Press: Stanford, Calif, 1993: 329-357.

Velleman, D.J. 2003. Narrative explanation, *The Philosophical Review*, 112: 1-25.

Walker, M. 2007. Superlongevity and utilitaranism, *Australasian Journal of Philosophy*, 85: 581-595.

Weindruch, R., and Sohal, R.S. 1997. Caloric intake and aging, *The New England Journal of Medicine*, 337: 986-94.

Williams, A. 1997. Intergenerational equity: an exploration of the 'fair innings' argument, *Health Economics*, 6: 117–132.

Williams, B. 1973. The Makropulos case: Reflections on the tedium of immortality, in *Problems of the Self*, Cambridge: Cambridge University Press, pp: 82-100.

Wisnewski, J. J. 2005. Is the immortal life worth living? *International Journal for Philosophy of Religion*, 58: 27-36.