

Note: this is an article whose final and definitive version is published in *Ethical Theory and Moral Practice*; the official online version of this article is available via: <http://link.springer.com/article/10.1007%2Fs10677-013-9431-6>.

Life in over-abundance: Agar on life-extension and the fear of death

Aveek Bhattacharya and Robert Mark Simpson

Abstract: In *Humanity's End: Why We Should Reject Radical Enhancement*, Nicholas Agar presents a novel argument against the prospect of radical life-extension. Agar's argument hinges on the claim that extended lifespans will result in people's lives being dominated by the fear of death. Here we examine this claim and the surrounding issues in Agar's discussion. We argue, firstly, that Agar's view rests on empirically dubious assumptions about human rationality and attitudes to risk, and secondly, that even if those assumptions are granted, the fears that Agar adverts to are unlikely to dominate people's lives if and when radical life-extension is made possible. Further, we claim that the structure of the decision-making process around life-extension is unlikely to be the way that it would have to be in order for Agar's claims about fear of death to make sense. Finally, we argue that Agar is implicitly committed to a narrow conception of human value. In response, we suggest that the pursuit of life-extension can itself be seen as an expression of certain important aspects of our distinctively human nature.

1. Introduction

Different evaluative perspectives can be adopted in considering the ethical implications of human enhancement technologies. We may assess things from the perspective of the individual who stands to benefit from the proposed enhancements, e.g. by asking whether technology x will truly benefit the individual, or will instead turn out to be a 'poisoned chalice', enhancing the individual's abilities in some

respects to the detriment of other aspects of his life. Alternatively, we may assess things from the perspective of people in positions of social influence, e.g. political leaders, legislators, and people with enormous wealth at their disposal. When adopting this perspective, the benefit or otherwise to individual users is only part of the story. Where technology x holds genuine benefits for individuals, someone in a position of influence might still think its development should be opposed if she judges that severe costs to others or gross distributional injustices will arise due to technology x. Conversely, where technology x seems likely to be a ‘poisoned chalice’ for those who have access to it, one may still think it right to support – or at least, to not oppose – x’s development, since one may judge that the damages involved are marginal, and that individuals should decide for themselves whether to risk incurring these damages. The ethical quandaries we encounter in navigating between these different perspectives on human enhancement are structurally similar to those that figure in perennial debates in political philosophy, e.g. about how states should balance individual liberties against the common good, and about the limits of paternalistic intervention to prevent individuals from making choices severely detrimental to their own welfare. This is, we suggest, precisely how ethical questions about human enhancement should be approached. It is no good thinking about the benefits and possible hidden costs of new technologies as if they were nothing more than a novel gamble that individuals may consider undertaking for themselves. Whether such gambles are likely to play out well or badly for the individual, we must also think about the situation of choice that *everyone* will find themselves inexorably drawn into once a given technology is realised, and whether that situation is something to promote, or instead to try to resist.

One class of potential human enhancement technology that calls for this dual-perspective ethical approach is life-extension technology. If it becomes possible to extend people’s lives far beyond their current duration, this will be a profound change in the conditions of human existence. While there are certainly great distributional injustices that could arise in the face of such technology, it is less plausible, *prima facie*, to suppose that life-extension would represent a poisoned chalice for those who have access to the relevant technology. The benefits of an increased healthy lifespan to individual users are plain to see. Attempts to identify harms to users caused by their increasing lifespans typically devolve into hazy and unpersuasive talk about the importance of ‘the given’, or about the moral hazards of hubris (see §2). In recent work, though, Nicholas Agar has suggested an entirely different kind of reason why we might

expect radically extended lives to be a poisoned chalice. In *Humanity's End: Why We Should Reject Radical Enhancement* (2010), Agar argues that radically extending our lives – or becoming ‘negligibly senescent’, to use a term favoured in some circles – will result in people’s lives being “completely dominated” by the fear of death (114).¹ Agar’s argument starts from the individual user’s perspective: he thinks it would be a bad thing for the individual, by her own lights, to opt for negligible senescence. Given this, one might wonder why Agar bothers to espouse his view. If it really *would* be bad for individuals (by their own lights) to undertake radical life-extension treatment, why not just trust people to not do what is injurious for them by their own lights? What animates Agar’s argument is the idea that while life-extension will be a poisoned chalice for individuals, it will nevertheless be an extremely tempting prospect once it is made possible. Thus, while Agar’s claims are primarily addressed to individual users, they also engage with a socially-oriented ethical perspective. Agar’s argument is motivated by his anticipation of the emergence of a destructive economy of desire and incentive where radical life-extension has become an option, and where individuals have been left to decide for themselves – as liberals like Agar prefer in principle² – whether they will pursue it. Indeed, the corrosive fear of death that is central to Agar’s argument may enter the frame not only once a person has actively undertaken to extend her life, but as soon as the option to significantly extend her life becomes available. Consequently, if Agar is right that the fear of death will spoil people’s lives in an era of negligible senescence, there are normative implications in this not only for individuals – of the ‘buyer beware’ type that Agar emphasises – but for legislatures, voters, funding bodies, and researchers, insofar as all of these can shape the potentially perilous situation we eventually find ourselves in regarding the prospect of radical life-extension.

Given that life-extension technologies are already being researched in earnest,³ it may come as a relief to find that Agar’s argument from fear of death, while it *is* a significant improvement on previous ethical arguments against radical life-extension, ultimately remains unconvincing. In §2 we briefly survey the recent ethical literature around radical life-extension. In §3 we summarise Agar’s argument, before

¹ Parenthetical page numbers in what follows all refer to *Humanity's End*. Agar discusses a variety of human enhancement possibilities in *Humanity's End*; here, however, we confine our attention to his arguments about life-extension.

² Agar espouses liberal principles in the development and use of human enhancement in his *Liberal Eugenics* (2004).

³ See for instance Bodnar et al. (1998), Ruan et al. (2002), Finch (2009), Baumer (2011), and Jaskelioff et al. (2011).

identifying and criticising some of the assumptions about human attitudes towards death, and about human rationality in general, upon which it rests. In §4 we consider a possible reply to these criticisms, namely, that individuals who undergo life-extension will become more preoccupied with the risk of death than are most of us living now, under conditions of ‘normal senescence’. In §5 we argue that even if we were to grant Agar’s assumptions about human attitudes towards the risk of death, not even rational utility-maximisers would suffer the consuming fear of death that he imagines. We also argue that the structure of the decision-making process around future life-extension technology need not be, and probably would not be, the way it would have to be in order for the prospect of ‘lives spoiled by fear of death’ to be intelligible. In §6 we argue that Agar is implicitly committed to a relatively narrow conception of human value, and in response, we identify alternative aspects of human nature and human value which we see as expressed in, rather than imperilled by, the pursuit of radical life-extension.

2. The debate around life-extension

Much of the literature on the ethics of life-extension simply deals in discussion and conjecture about the future prospects of the relevant technologies. At the modest end of the spectrum there are authors who consider the possibility of typical lifespans and health-spans⁴ increasing to approximately double their current length (Singer 1991; Glannon 2002a; Walker 2007). At the more speculative end of the spectrum, others advert to the possibility of lives extended towards 1000 years, or even extended *indefinitely*, i.e. not to the point of immortality as such, but as long as possible within the constraints imposed by (i) finite material resources, (ii) fatal accidents and (iii) the finite lifespan of the universe.⁵ An assortment of

⁴ Roughly, the health-span is the period during which a satisfying level of vitality and ability are maintained, whereas the lifespan is the period of ‘mere’ bodily survival. Not even enthusiastic advocates of life-extension express any interest in extending the lifespan without also extending the health-span. (The perils of such a scenario, writ large, are discussed in Fukuyama 2002.) In what follows we will not labour the distinction. When we’re talking about radical life-extension, negligible senescence, etc., we will assume throughout that what stands to be extended is health-span rather than lifespan.

⁵ Aubrey de Grey is the most notable voice in this camp; he sometimes describes life-extension research as an attempt to find “the cure for aging” (de Grey 2005: 663). De Grey heads up the Strategies for Engineered Negligible Senescence Foundation, a

technological advances have been identified as potential routes towards dramatically extended human life.⁶ At the genetic level, the hope is that it will become possible to manipulate the telomere nucleotide sequences at the ends of chromosomes, which regulate chromosome replication, and whose normal deterioration over time induces cellular senescence. The reversal of aging effects in mice has already been achieved via this route (Jaskelioff et al. 2011), and together with other novel advances – e.g. in hormone therapy, or caloric restriction, or synthetically produced bodily organs – this development raises the possibility of a dramatically extended human lifespan (see Marshall 2006; Barazetti 2011). We do not pretend to have the expertise required to assess the likelihood of the predictions we find in the literature. Here we will follow Agar in his moderate credulity towards those who expect to see significant life-extension (e.g. a doubling of the human lifespan) at some point in the 21st century. Like Agar, and others, our interest is in why such an expectation may occasion major ethical reservations.

The argumentative starting point for advocates of life-extension technology is often simply to explain why, for most of us, it would be a good thing to be able to extend our lives – why this is the kind of thing that many of us would want, and why we would be reasonable in wanting it. Occasionally this part of the discourse takes a critical bent, responding, for instance, to the suggestion (commonly linked to Bernard Williams 1973), that a significantly longer life may be excruciatingly boring. Much of the time, though, it simply amounts to an assurance that more life will mean more time to enjoy valuable things: friendship, love, creativity, satisfying work, aesthetic pleasures, and so on. So far as reassurance is needed, this may be due to the worry that in altering the conditions of our senescence, we will alter our capacity to find value in such things, or to find meaning of the same important type in our valuations of such things. John Harris offers an assured and straightforward response to these kinds of worries. In having her sight restored, he says, the blind person may alter some important or essential part of who she is, and this may

research and advocacy group based in California, whose aim is to “ensure widespread access to rejuvenation biotechnologies which comprehensively address the disabilities and diseases of aging” (see www.sens.org).

⁶ The literature typically focuses on extended lifespans in bodies functionally similar to those we currently have, rather than the prospect of people ‘uploading’ their minds into cyberspace. Agar does also examine ‘uploading’ independently of his discussion of life-extension in *Humanity’s End*. He argues that uploading one’s mind is irrational if there is a non-zero probability that this would constitute one’s death, e.g. if the substrate of one’s ‘mind’ would, post-uploading, merely be *simulating* feelings, memories, etc. This argument has generated some discussion in the artificial intelligence literature (see Levy 2011; Agar forthcoming).

profoundly affect the character of her valuations. However, so Harris maintains, this would not render her choice to have her sight restored irrational. Her motive remains intelligible to her, and to us, even if there remain lingering doubts or puzzles about how her ‘valuing self’ stands in relation to the goods she enjoys, post-operation; and the same goes for life-extension (Harris 2004: 531). Another approach in explicating the ‘commonsensical self-interest’ view of life-extension is to highlight what Aubrey de Grey calls “the uncontroversial repugnance of denying others the choice to live” (2005: 661). If we can prolong the life of a healthy person who wants to live, de Grey argues, then it is (*pro tanto*, let’s say) wrong to act, or to perform an omission, so as to deny this possibility. Continued life for those who desire it is something that we naturally regard as a fundamental moral entitlement, and so if we can extend this entitlement beyond the limitations that we currently face, then we should, de Grey argues, devote resources – or at the very least, we should not *object* – to the pursuit of that end.

On the opposing side of the debate there is a mixture of firmly ‘bio-conservative’ authors, such as Leon Kass (2001, 2003) and Michael Sandel (2004), and other more liberal-leaning authors (by reputation, at least), such as Jürgen Habermas (2003) and Bill McKibben (2003). A central theme in the work of these authors is the moral importance of life’s ‘given-ness’. The idea, in broad strokes, is that there is great value in our living under and making choices within the constraints that our humanity imposes upon us, rather than exercising choice over the basic parameters of human existence. There is a spirit of humility behind this notion, the idea being that a good life consists in diligently striving after value with the powers one has, rather than grasping for further powers with which more and better value might, so one hopes, be attainable. We find this contrast dubious, though. To explain the value of choice and effort being constrained by ‘the given’, it seems we have to impute some moral specialness to choice and effort. But then it is unclear why choice and effort should lose this moral specialness when they are exerted or exercised over the parameters of our lives (on this line of argument see for instance Horrobin 2006: 282-86). A simpler conservative argument simply calls into question the wisdom of ushering in dramatic, unpredictable changes in our reproductive lives and vocational structures, of the sort that would

seemingly be inevitable given the possibility of radical life-extension, and which might well, for all we can surmise, leave society worse off (see Callahan 1994; Hackler 2004; Temkin 2011).⁷

Agar's arguments are, as we previously noted, addressed to the individual's perspective; he aims to show, contra de Grey, Harris, and others, that negligible senescence would be worse *for the individual*, apart from any social problems that it might produce. It is a meritorious feature of Agar's work, we think, that it aims to bridge the communicative impasse between opponents of life-extension who speak of 'the importance of the given', and the 'threat to our humanity', and those who, favouring the ethical language of rights or autonomy, find such terms somewhat opaque. Granted, it is not evidently *false* that our humanity will be threatened by the realisation of radical life-extension; we may well be ushered us into some form of post-human life or society. But the question, as ever when this point is raised, is why conservative ethical conclusions should follow from this judgement (see for instance Allen Buchanan's *Beyond Humanity?*). It is easy to form the suspicion that appeals to humanity are just puffed-up expressions of status quo bias (see Bostrom and Ord 2006), or expressions of personal preferences ('I don't want to live in a world like that') repackaged in ethical language. Agar's work makes progress here by offering a more perspicuous account of what it is that life-extension technology (allegedly) threatens. His claim is

⁷ Some authors like David Gems (2003) have a foot in both camps; he agrees with de Grey, Harris, et al. that life-extension would benefit the individual, but he sees the likely social consequences of life-extension technology, in particular the potential for power to become more concentrated over time, as a decisive reason against its development. Alongside the debate between advocates of life-extension and conservatives, there are numerous pieces in the literature examining how that debate is framed and what implications different ethical theories have for it. Some authors, concerned with Parfit's (1984) distinction between mere *survival* and identity-preservation, ask whether we can expect radically-extended human lives to allow for identity-preservation across time, or instead for mere survival, or for some lesser continuity relation, and what the ethical consequences would be given each prospect (see Glannon 2002b; Harris 2002a, 2002b; Schloendorn 2006; Barazetti and Reichlin 2011). Others draw on economics research into the trajectory of people's happiness levels across the lifespan, to formulate predictions about what might be at stake, from a purely utilitarian perspective, in prolonging human lives (Walker 2007; Blackford 2009). Others examine parochial judgements – e.g. about elderliness, gender, and generational change – which lead to a moralisation of ideas about what is 'normal' in human society, and the consequences of this in life-extension (Overall 2003). And finally, some authors question whether bioethics as a discipline is adequately-equipped to address ethical questions about life-extension, given that the kind of anxieties that life-extension generates cannot be readily articulated in bioethics' prevailing idiom (Trotter 2004).

that our understanding of, and attitudes towards, our own mortality, however uneasy they may be at present, would become worse still if and when radical life-extension becomes possible.

3. The argument from fear of death

As we said in §1, Agar's principal argument against life-extension stems from his suspicion that "the fear of death may completely dominate the lives of negligibly senescent people" (114). He observes that a crucial difference between negligible senescence and immortality is that in the former case, the possibility of death remains. So, while aging and disease may be conquered, or reduced to remote dangers, death from non-biological causes, most significantly through accident and injury, will remain a threat. As accident and injury become the primary causes of death for the negligibly senescent, and as negligible senescence massively increases the potential cost of accidents, Agar anticipates that people will become increasingly cautious. He quotes de Grey as suggesting that "once we cure aging, driving (even on the ground!) will be outlawed as too dangerous for others" (115). Currently the victim of a fatal accident stands to lose fewer than eighty years of life. However, in an era of negligible senescence, the accident-victim may stand to lose centuries of life. Consequently, so Agar and de Grey postulate, people who have undergone life-extension treatment will avoid even very remotely risky activities. Agar illustrates the envisaged situation with an example of someone deciding whether or not to drive to the cinema. At present, most of us face this decision expecting a high probability of a brief, enjoyable experience (i.e. seeing the film) and a low probability of an exceptionally bad outcome, e.g. death in a car accident. For most of us, the expected benefit clearly outweighs the potential cost. For a negligibly senescent person, however, the benefit remains modest, while the potential cost is massively increased. No movie will be good enough to outweigh the prospect of an early death, and so the negligibly senescent will choose to stay at home. This rationale applies not only to driving to the cinema, but to all sorts of everyday activities. The negligibly senescent would become paralysed by fear, house-bound, and restricted to a simulation of outdoor activities in virtual worlds. Radical life-extension, Agar says, will

alienate us from the things and people who currently give meaning to our lives... Currently, we get pleasure from direct, unmediated connections with parts of our environments. We swim in the

sea, climb mountains, and ski down snowy slopes. Part of the pleasure we derive comes from the fact that they take us slightly out of our comfort zones – they seem a little bit dangerous... The little bit of danger that makes snorkelling tropical reefs or flying a micro-light aircraft exciting for us is likely to translate into insane recklessness for negligibly senescent people... I predict that negligibly senescent people will retreat from the world.” (Agar 2010: 122-23)

The first thing to notice is that this prediction seems to assume that people’s attitudes and behaviour roughly follow the canons of rational choice theory. Indeed, at several points Agar couches his discussion in the language of expected utility theory. He suggests that people make risky decisions in a way that is sensitive to how potential costs and benefits weigh against one another, scaled according to the likelihood of their occurrence. This is evident, for instance, when he speaks of someone driving to the cinema as facing “a very high probability of a brief but enjoyable experience – seeing a movie – minus the miniscule probability of a very bad outcome – death” (116). Now, it is well-established that people do not typically deliberate in explicitly probabilistic terms, let alone apply the formal devices of rational choice theory, when facing risky decisions (e.g. see Huber et al. 1997). But then Agar does not and need not claim that people explicitly engage in expected-utility calculations when faced with risky decisions. What he needs to claim is that in our decision-making about actions which carry a remote mortality-risk, we are typically responsive to probability-weighted costs and benefits in something like the way they are accounted for in the apparatus of rational choice theory. While Agar concedes that “humans are not entirely rational in our responses to risk” (115), his argument requires that where the risk of death is concerned, our choices will at least be *sensitive* to the anxieties that would be felt by rational utility-maximisers.

The principal problem with Agar’s argument, then, is that this simply isn’t how remote risks of death are confronted in human attitudes and behaviours. If people were typically sensitive to expected-disutility considerations, then we would become gradually less averse to life-endangering activities as we grow older and have fewer years of life to lose in the event of an accident or misfortune; or as Agar says, “the young should be the ones watching *Days of Our Lives* reruns in the safety of their lounges” (116). But generally speaking the opposite is true. The evidence, both from surveys and from studies in behavioural psychology, bears out the familiar notion that people are progressively less willing to take risks as they grow older (Dohmen et al. 2005). This may be because people simply don’t consciously deliberate about

risks, or it may be due to older people being less prone to optimism, and thus less likely to underestimate the size of risks, as some studies have indicated (DeJoy 1989). One way or the other, the way people face up to mortality-risk across the course of their lives undermines any notion that our decision-making in this area is sensitive to probability-weighted costs and benefits.

Moreover, this tendency towards irrational optimism presents a further problem for Agar's argument. Studies have demonstrated that people see negative events (of any sort) as less likely to happen to them than to others (e.g. see Weinstein 1980). Whatever the dangers of a particular activity, people typically find themselves guided by a sense that 'it won't happen to me'. This optimism is especially strong when people believe they have control over the outcome in question, since people also tend to have excessive faith in their own abilities. Studies have found, for instance, that a large majority of drivers believe themselves to be safer and more skilful than average (e.g. Svenson 1981). Frequently, then, when we do recognise remote dangers in the activities we undertake, we will irrationally see ourselves as immune from those dangers.

Even if we can be made to recognise the remote dangers we face in our everyday activities, human minds are notoriously bad at incorporating this sort of data into decision-making processes. As Kunreuther, Novemsky and Kahneman explain, "people given low probabilities associated with some event, may not know how to evaluate these likelihoods. It is difficult to gauge how concerned one should feel about a 1 in 100,000 probability of death without some comparison points. Most people just do not know whether 1 in 100,000 is a large risk or a small risk" (Kunreuther et al. 2001). The problem for Agar, then, is not merely that people typically have distorted expectations about the likelihood of their succumbing to a very remote risk of death, but also that people struggle to make sense of how they ought to behave, from a prudentially self-interested perspective, even when they have been disabused of their irrational optimism.

The upshot of these things is that human beings for the most part pursue everyday actions and decisions in a manner that simply *ignores* very remote risks, including very remote risks of dying. With respect to an everyday activity like driving, for instance, Camerer and Kunreuther (1989: 570) observe that the chance of a fatality for an average driver is "well below the threshold many people set before paying attention to a risk". Stated in expected utility terms, the idea is that human beings characteristically treat

driving as an activity which – at least as far as the prospect of one’s own death goes – has *zero* expected disutility. Obviously there is a sense in which this is a plainly incorrect appraisal of one’s situation. But then it is not hard to see why, even upon reflective scrutiny, we might welcome our habit of facing up to remote dangers in this fashion. As Slovic et al. say (1977: 254), “there are only so many things in life one can worry about”. And “without some sort of threshold for concern, people would spend their entire lives obsessively protecting themselves against a ‘Pandora’s urn’ of rare horrors”. Their suggestion, then, is that our sub-rationality in this arena is precisely what allows us to avoid the scenario that Agar foresees, i.e. our becoming preoccupied with the innumerable remote risks that we face. Agar’s claim, that people who undergo life-extension will have their lives spoiled by fear of death, only makes sense if people are generally attuned to, and responsive to, rational risk-assessments when faced with very-low-probability risks. The fact that human beings are *not* typically like this when it comes to the risk of death is what enables us to go about our lives without constantly worrying about very unlikely tragedies befalling us, and what will enable people in the future to do likewise in the face of radical life-extension.

4. Negligible senescence and rationality

Our claim is that Agar’s argument is unpersuasive because of ordinary human beings’ normal sub-rational attitudes, specifically, our disinclination to attend to, and difficulty in fathoming, the mortality risks associated with everyday activities. Would idealised utility-maximisers be preoccupied by the fear of death upon attaining negligible senescence? Maybe. But negligibly senescent *humans* – just like ‘normally senescent’ humans – will generally downplay or just ignore remote mortality risks.

How might Agar respond to this? One possibility to consider is that people who undergo life-extension treatment will have different attitudes to typical human beings today, in a way that makes them more susceptible to becoming preoccupied by fear of death. This is what Agar seems to have in mind when he says that “de Grey’s disciples are likely to have attitudes to risk that are more rational than those of teenage boy drag-racers” (116). If negligibly senescent people are like ideal utility-maximisers in the way they think about mortality risks, perhaps they will be consumed by fear of death.

But is there any reason to suppose that negligibly senescent people will actually be this way? Perhaps people living in an era of negligible senescence will be technologically aided in their ability to calculate expected utilities for low-probability risks. Even if we allow the conjecture, though, this is a red herring. Anyone equipped with the internet and a calculator should already be able, with a little effort, to usefully estimate the risks associated with low-risk activities, like riding one's bicycle into town, or going for a walk in the hills, or attending an air show. Our sub-optimal rationality in this area is less about the ability to acquire information and calculate probabilities, and more about our preparedness to pay careful attention to, and be guided by, judgements about very remote risks of death. Instead, it may be that Agar thinks a selection effect is likely to occur, so that people who seek to undergo life-extension treatment are from the outset more attentive to mortality risks than most normally senescent people, and thus more prone to being afflicted by fear of death. It is not too far-fetched to imagine that people who eagerly embrace life-extension developments will be more inclined to apply expected-utility-type calculations to their everyday decisions about risky activities. (For instance, it seems that de Grey himself, at least, thinks about these matters in the way that Agar describes, and is accordingly very concerned about how negligibly senescent people can negotiate mortality anxieties.⁸) On the other hand, though, it seems unlikely that any such selection effect would have wider-reaching negative consequences with respect to life-extension treatment. People like de Grey, who are already preoccupied with mortality and how it might be delayed, will presumably remain thus preoccupied (albeit with raised stakes) if and when negligible senescence becomes possible. But the majority who do not obsess over their deaths can surely avoid being drawn into this: they can either forego life-extension technology, or use it while retaining the sort of mindset that most of us have at present, under which routine and remote risks of death are for the most part ignored. Alternatively, Agar might claim that the experience of life-extension treatment *itself* will make the negligibly senescent more conscious of their mortality, and therefore more fearful about the ubiquitous but remote mortality-risks they face. Undergoing life-extension raises the stakes of death, in terms of how much life each individual stands to lose, and it also represents a conscious choice to try to delay one's death, of the sort that might place mortality at the forefront of the individual's thoughts. But

⁸ Hence de Grey's comment about banning driving (see §3); de Grey has also suggested, in personal communication with one of the authors, that the argument from fear is among his greatest worries about negligible senescence.

this is mere conjecture, and it is not borne out by observable patterns in our attitudes at present. That is to say, neither highly health-conscious people (i.e. those who consciously act in the interest of longer-living), nor young people (i.e. those who have more years life to lose) seem to be any more neurotic about remote risks of death than the rest of the population. In the absence of anything more reliable to be guided by, our predictions about the attitudes and cognitive dispositions of negligibly senescent humans should be based on what we can observe in human beings today. And under current conditions, most of us find ourselves with a felicitously imperfect way of confronting remote risks of death. On one hand, we are inattentive to the matter, and poorly-equipped for making precise expected-utility calculations about the risks involved. But on the other hand, these things spare us the angst of worrying about death at every turn. Without a compelling reason to think these patterns will be infelicitously altered following the realisation of negligible senescence, Agar's argument has little purchase. Some negligibly senescent people may obsessively fret about their mortality, but this is unlikely to be a worse or more widespread condition than in the few normally senescent people obsessively fretting about *their* mortality today.⁹

5. Rational fears and vexed decisions

In arguing that radical life-extension will not lead to people being consumed by fear of death – and thus deterred from engaging in everyday forms of leisure – we have granted, thus far, that radical life-extension may well lead an *ideal utility-maximiser* to be haunted by fear of death in the way Agar imagines. But this is

⁹ Some studies find that while people are generally hyperbolic discounters of future rewards, we conform to the exponential discounting patterns of classical economics more the longer we live (see e.g. Green et al. 1994); and one might argue that this is *pro tanto* evidence that negligibly senescent people will approximate the attitudes of rational utility maximisers. The argument, roughly, would be that since long-living (normally senescent) humans approximate the attitudes of rational utility-maximisers in one respect (discounting of future benefits), we can expect *very* long-living (negligibly senescent) humans to approximate the attitudes of rational utility-maximisers in other respects, including appraisal of mortality risks. But this is a dubious inference. The best explanation of the data on the relation between discounting benefits and age is that approaching the end of one's life makes one attend more carefully to the value of future benefits. Proximity to life's end, rather than living longer *per se*, drives the attitudinal change. Thus, although negligible senescence increases years of life, it actually *delays* the key factor (i.e. proximity to death) which most plausibly explains the closer approximation of rational utility-maximisation in older (normally senescent) people. We are grateful to an anonymous reviewer from this journal for bringing this point to our attention.

dubious in its own right. Agar says that “for the car trip to the cinema to be as uncontroversial a matter for the negligibly senescent person as us, the experience of seeing the movie would have to be over one hundred times better than it is for us” (117). His idea is that if person A has a lifespan x times greater than ours, then she suffers disutility in death x times more than we will, and so movies will have to be x times more enjoyable for A than for us in order for A to go to the movies with the same peace of mind that we do. First, a minor point: this line of reasoning assumes that the frequency of transit accidents will remain roughly constant across the two scenarios. It is likely, though, that fatal accidents will continue to decline in frequency into the future. (The rate of death per mile driven in the US is five times lower today than it was in 1950; see Levitt and Dubner 2009: 146.) In any case, we can suppose for the sake of argument that life-expectancies increase by a factor greater than or equal to future decreases in accident frequency. The question then, is why we should also suppose that improvements in the cinematic experience would have to fully keep pace with increases in lifespan in order for it to still be rational to go out to the cinema. Agar’s claim only makes sense if we suppose that for normally senescent people nowadays, going to the cinema, and other remotely risky leisure activities, are only *barely* rational options from a utility-maximising perspective. And this is surely false: for most of us, if the question were ever to arise, the expected benefits of our everyday leisure activities would overwhelmingly outweigh the costs associated with remote risks of death. Indeed, Agar tacitly acknowledges this when he talks about the great value that resides in everyday leisure, whether of the mundane kind, or the more thrill-seeking kind. For the negligibly senescent, a great deal stands to be lost in the event of a fatal accident, but – for them, like us – a great deal *also* stands to be lost in retreating from all risky activity.

This leads us to a more general problem with Agar’s argumentative approach. His worries about a consuming fear of death only come into view when we are thinking from within a utility-maximising perspective, about how to balance (a) the value of very-slightly-risky activities, against (b) the value of a very long but cautious and risk-averse life. When things are framed things in this way, though, we just don’t have to worry about individual utility-maximisers getting things wrong from a utility-maximisation standpoint. Subjective tastes and preferences will influence each individual who encounters the choice, and each individual can choose whichever package of costs and benefits is preferable to her. Those, like Agar, who would rather risk their necks occasionally than live a long life with all the fun and dangerous

parts taken out, can do just that. Those who greatly value a very long life, even if it means foregoing a wide range of risky activities, can opt for that. Of course, as we stressed in §4, human beings are not ideal utility-maximisers. Our point is that in order to make sense of why negligibly senescent humans might retreat from the world, as in Agar's prediction, we must treat utility-maximising as a near-enough approximation of how they will respond to risk. And this renders Agar's argument self-undermining, because to whatever extent people's attitudes *do* approximate those of rational utility-maximisers, it makes no sense to see people as being made worse off by having a new option in the form of life-extension. Either the option is an agreeable one and they pursue it, or it isn't and they don't.

One worry may be that the choice to extend one's life will often be made impetuously, before individuals have made a sober appraisal of the costs that negligible senescence may carry. It is one of Agar's themes in *Humanity's End* that the way human enhancement research and technology is marketed – to investors, governments, and the general public – is skewed in order to emphasise potential benefits and downplay potential costs. His concerns on this front are entirely reasonable, and all parties should be able to agree that future consumers of enhancement technologies must be supplied with information, from diverse sources, about the costs and benefits that may be at stake in their choosing to undergo treatment (including information about disagreements concerning those putative costs and benefits). Obviously errors can still occur given these sorts of precautions. For instance, someone who shares Agar's tastes and preferences may not realise the vexing dilemmas she will face having undertaken life-extension treatment. But in the case of life-extension, at least, such mistakes need not be costly, since they are easily reversible. Extending one's life does not mean becoming immortal! Moreover, people who want to opt *out* of life-extension treatment after having opted in would almost certainly not face the dreadful prospect of taking their own lives. As Mark Walker (2007) points out, any remotely imaginable life-extension technologies would involve ongoing treatments at regular intervals, as opposed to a one-off, 'have-this-operation-and-enjoy-the-next millennium', transformation. The beneficiary of life-extension treatment who finds herself being troubled by unbearable anxieties about the risk of a fatal accident need only cease her treatments, and ease back into the less-anxious mind-set of normal senescence. We doubt Agar's predictions about

how people will react to life-extension, but even if (for some) life in over-abundance *is* ruined by fear of death, people will surely be able opt out of the condition that causes the anxiety.¹⁰

Again, therefore, it seems to us that Agar's objection to the prospect of life-extension treatment cannot ultimately be that people will become worse-off by making the wrong choice about whether to pursue such treatment. His underlying concern, rather, is that we will be made worse off simply by virtue of having the option to pursue such treatment. This notion – that the mere fact of having a choice can sometimes make the individual worse off – is a familiar one in bioethics; we see it for instance in David Velleman's claim that giving people the option of undergoing assisted suicide will make some people worse off, even by their own lights (see for instance Velleman 1992). In the present case, we know that there are already people, like de Grey, who are preoccupied with mortality and with how it might be resisted. Agar is concerned that the rest of us will find ourselves being drawn into this sort of anxious mindset, just by having to confront the question of whether or not to undergo life-extension treatment. As we argued in §4, this kind of speculation is unsupported by observations of current patterns in the way that people face up to remote the mortality risks they face. Further to that point, though, we think Agar's conjecture reflects an unduly pessimistic view about how it is that people find themselves afraid of death and/or neurotically preoccupied with their own mortality. It is surely no accident that the typical way in which we find value in our activities and ways of life is to appreciate the positive states that we value, firstly, while only secondarily worrying about the dangers that might deprive us of further appreciation. When viewing a beautiful vista, our instinctive response is to admire its beauty, not to worry about the possibility of one day becoming vision-impaired. When it comes to valuing life itself, some may find this order reversed, but most of us just appreciate the good things, without worrying about their transience. Unlike Agar, we think such dispositions are robust enough to endure in the face of life-extension technology. If the way we already experience valuable activities in life is any guide, negligible senescence will be greeted as an opportunity for more value, and not as a paradoxical threat to value.

¹⁰ Observe John Harris's response to the suggestion that negligible senescence would turn out to be a poisoned chalice: "I would (as of now) be quite happy to sample a few million years and see how it goes" (Harris 2002b: 284).

6. Human nature and value

It is important to observe, by way of conclusion, a significant symbolic divide that manifests itself in Agar's discussion (and to some degree in ours as well), between those who eagerly pursue the prospect of life-extension – who Agar sees as destined to “spend a great deal of time on the internet” (120) – and the regular folk who enjoy “sailing, bush walking, and wind surfing” (122), and who Agar encourages to join him in puzzling about the peculiar people who would renounce such things for the sake of a longer life. This symbolic divide is starkly revealed, for instance, when Agar says that negligibly senescent people may still have “worthwhile and rewarding lives”, even though “*We* may find the excitement of [their] activities difficult to grasp” (127, our emphasis). Now, although we think Agar is wrong to expect a preoccupation with risk-avoidance to accompany radical life-extension, we do largely agree with him that this kind of preoccupation alienates people from some of the most valuable things in life. A person who would stay indoors in a death-proof bunker indefinitely in order to have a very long, risk-free existence, is surely missing out on some crucially important part of what it means to be human. At the same time, though, we think Agar is too quick to see advocates of life-extension as people that have become alienated from recognisably human goods altogether. There needn't be anything *un-human* or *anti-human* driving the *de Greys* of this world to pursue life-extension. The human species may, after all, be distinguished from other parts of the animal kingdom in terms of our instinctive compulsion to try to technologically modify the conditions of our own existence. From this perspective, those who seek to make it possible for us to live much longer lives can be understood as acting out an instinct or disposition which is as essentially human as any. When we are thinking about the potential costs and benefits of life-extension for the individual, then, we should not think of the divide as being between reckless technophiles and faithful custodians of human value. Rather, the conflict here is over which *aspect* of our humanity will be given priority: the part that is content making the best of our lives as we find them, or the part that is compelled to transform the fundamental parameters within which our lives are led. And there is surely something to be said on both sides, as far as the costs and benefits for the individual are concerned.

Radical life-extension technology has the potential to drastically alter the shape of individuals' lives and of society as a whole. When considering such drastic changes, it is obviously reasonable and prudent to think hard about whether the superficial benefits at stake will be offset by hidden costs. Agar's

approach to this matter is original and thoughtful. Where others employ mystifying language to express their ethical reservations about life-extension, Agar clearly identifies a state of affairs that would be bad for those who experience it, along with a *prima facie* plausible explanation as to why that state of affairs might come to pass given the possibility of radical life-extension. It is absolutely correct, after all, that the negligibly senescent *would* have many more years of life to lose than we do today. Contrary to Agar's conjectures, however, this will not result in people's lives being dominated by the fear of death. We can reasonably expect that people living in an era of negligible senescence will ignore and downplay the remote dangers they face in their everyday activities and recreational pursuits, in the same way that people today typically do. There are no good reasons to think that the characteristic attitudes and cognitive dispositions that people currently betray in confronting remote dangers in their lives will be any different if radical life-extension is made possible. Nevertheless, it is essential that we face up to the concerns Agar adverts to, and assure ourselves that life-extension *is* unlikely to be a poisoned chalice. If we are treating life-extension as a potential mixed blessing for those who pursue it, we can all too easily lose sight of the new and extreme inequalities that may arise in a world where the extremely powerful and wealthy can cement their power across multiple normally senescent human lifespans. The first-blush appraisal of life-extension technology – that those who have access to it stand to benefit enormously – holds up in the face of critical scrutiny. But this merely *clarifies* the ethical quandaries surrounding life-extension, rather than dispensing with them. The key question is not whether life-extension would be ruinous for those with the chance to pursue it. The question is whether it would be ruinous for everyone else.

Acknowledgements We are grateful to Simon Rippon, Nicholas Agar, and three anonymous referees from this journal, for their thoughtful feedback and criticism in response to this paper.

References

- Agar, Nicholas (2004), *Liberal Eugenics: In Defence of Human Enhancement* (Oxford: Blackwell).
- _____ (2010), *Humanity's End: Why We Should Reject Radical Enhancement* (Cambridge Massachusetts: MIT Press).
- _____ (forthcoming), "On the Irrationality of Mind-Uploading: A Reply to Neil Levy", *AI & Society*.

- Barazetti, Gaia (2011), "Looking for the Fountain of Youth: Scientific, Ethical, and Social Issues in the Extension of Human Lifespan", in Julian Savulescu, Ruud ter Meulen, and Guy Kahane (eds.), *Enhancing Human Capacities* (Oxford: Wiley-Blackwell).
- Barazetti, Gaia and Massimo Reichlin (2011), "Life-extension and Personal Identity", in Julian Savulescu, Ruud ter Meulen, and Guy Kahane (eds.), *Enhancing Human Capacities* (Oxford: Wiley-Blackwell).
- Baumer, Yvonne, Beate Scholz, Svetlana Ivanov, and Burkhard Schlosshauer (2011), "Telomerase-Based Immortalization Modifies the Angiogenic/Inflammatory Responses of Human Coronary Artery Endothelial Cells", *Experimental Biology and Medicine*, 236 (6), 692-700.
- Blackford, Russell (2009), "Moral Pluralism Versus the Total View: Why Singer Is Wrong About Radical Life-extension", *Journal of Medical Ethics*, 35 (12), 747-52.
- Bodnar, Andrea G., Michel Ouellette, Maria Frolkis, Shawn E. Holt, Choy-Pik Chiu, Gregg B. Morin, Calvin B. Harley, Jerry W. Shay, Serge Lichtsteiner, and Woodring E. Wright (1998), "Extension of Lifespan by Introduction of Telomerase into Normal Human Cells", *Science*, 279 (5349), 349-52.
- Bostrom, Nick and Toby Ord (2006), "The Reversal Test: Eliminating Status Quo Bias in Applied Ethics", *Ethics*, 116 (4), 656-79.
- Buchanan, Allen (2011), *Beyond Humanity? The Ethics of Biomedical Enhancement* (Oxford: Oxford University Press).
- Callahan, Daniel (1994), "Manipulating Human Life: Is There No End in It?", in Robert H. Blank and Andrea L. Bonnicksen (eds.), *Medicine Unbound: The Human Body and the Limits of Medical Intervention* (New York: Columbia University Press).
- Camerer, Colin F. and Howard Kunreuther (1989), "Decision Processes for Low Probability Events: Policy Implications", *Journal of Policy Analysis and Management*, 8 (4), 565-92.
- de Grey, Aubrey D. N. J. (2005), "Life-extension, Human Rights, and the Rational Refinement of Repugnance", *Journal of Medical Ethics*, 31 (11), 659-63.
- DeJoy, David M. (1989), "The Optimism Bias and Traffic Accident Risk Perception", *Accident Analysis & Prevention*, 21 (4), 333-40.
- Dohmen, Thomas J., Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner (2005), 'Iza Discussion Paper No. 1730: Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey', (Bonn: Institute for the Study of Labor).
- Finch, Caleb E. (2009), "Update on Slow Aging and Negligible Senescence: A Mini-Review", *Gerontology*, 55 (3), 307-13.
- Fukuyama, Francis (2002), *Our Posthuman Future: Consequences of the Biotechnology Revolution* (New York: Picador).

- Gems, David (2003), "Is More Life Always Better? The New Biology of Aging and the Meaning of Life", *The Hastings Center Report*, 33 (4), 31-39.
- Glannon, Walter (2002a), "Extending the Human Life Span", *Journal of Medicine and Philosophy*, 27 (3), 339-54.
- _____ (2002b), "Identity, Prudential Concern, and Extended Lives", *Bioethics*, 16 (3), 266-83.
- Green, Leonard, Astrid Fry, and Joel Myerson (1994), "Discounting of Delayed Rewards: A Lifespan Comparison", *Psychological Science*, 5 (1), 33-36.
- Habermas, Jürgen (2003), *The Future of Human Nature* (Cambridge: Polity Press).
- Hackler, Chris (2004), "Extending the Life Span: Mythic Desires and Modern Dangers", *HEC Forum*, 16 (3), 182-96.
- Harris, John (2002a), "Intimations of Immortality: The Ethics and Justice of Life Extending Therapies", in Michael Freeman (ed.), *Current Legal Problems* (Oxford: Oxford University Press).
- _____ (2002b), "A Response to Walter Glannon", *Bioethics*, 16 (3), 284-91.
- _____ (2004), "Immortal Ethics", *Annals of the New York Academy of Science*, 1019 (1), 527-34.
- Horrobin, Steven (2006), "Immortality, Human Nature, the Value of Life and the Value of Life-extension", *Bioethics*, 20 (6), 279-92.
- _____ (2011), "The Value of Life-extension to Persons as Conatively Driven Processes", in Julian Savulescu, Ruud ter Meulen, and Guy Kahane (eds.), *Enhancing Human Capacities* (Oxford: Wiley-Blackwell).
- Huber, Oswald, Roman Wider, and Odilo W. Huber (1997), "Active Information Search and Complete Information Presentation in Naturalistic Risky Decision Tasks", *Acta Psychologica*, 95 (1), 15-29.
- Jaskelioff, Mariela, Florian L. Muller, Ji-Hye Paik, Emily Thomas, Shan Jiang, Andrew C. Adams, Ergun Sahin, Maria Kost-Alimova, Alexei Protopopov, Juan Cadinanos, James W. Horner, Eleftheria Maratos-Flier, and Ronald A. DePinho (2011), "Telomerase Reactivation Reverses Tissue Degeneration in Aged Telomerase-Deficient Mice", *Nature*, 469 (6 January 2011), 102-06.
- Kass, Leon R. (2001), "L'chaim and Its Limits: Why Not Immortality?", *First Things*, (May 2001), 17-24.
- _____ (2003), "Ageless Bodies, Happy Souls: Biotechnology and the Pursuit of Perfection", *New Atlantis*, Spring (1), 9-28.
- Kunreuther, Howard, Nathan Nomevsky, and Daniel Kahneman (2001), "Making Low Probabilities Useful", *Journal of Risk and Uncertainty*, 23 (2), 103-20.
- Levitt, Steven D. and Stephen J. Dubner (2009), *Superfreakonomics* (London: Allen Lane).
- Levy, Neil (2011), "Searle's Wager", *AI & Society*, 26 (4), 363-69.
- Marshall, Jennifer (2006), "Life-extension Research: An Analysis of Contemporary Biological Theories", *Medicine, Health Care and Philosophy*, 9 (1), 87-96.

- McKibben, Bill (2003), *Enough: Genetic Engineering and the End of Human Nature* (London: Bloomsbury).
- Overall, Christine (2003), *Aging, Death, and Human Longevity: A Philosophical Inquiry* (Berkeley: University of California Press).
- Parfit, Derek (1984), *Reasons and Person* (Oxford: Clarendon Press).
- Ruan, Hongyu, Xiang Dong Tang, M. L. Chen, M. A. Joiner, Guangrong Sun, Nathan Brot, Herbert Weissbach, Stephen H. Heinemann, Linda Iverson, Chun-Fang Wu, and Toshinori Hoshi (2002), "High-Quality Life-extension by the Enzyme Peptide Methionine Sulfoxide Reductase", *Proceedings of the National Academy of Sciences of the United States of America*, 99 (5), 2748-53.
- Sandel, Michael J. (2004), "The Case against Perfection", *The Atlantic Monthly*, (April 2004), 51-62.
- Schloendorn, John (2006), "Making the Case for Human Life-extension: Personal Arguments", *Bioethics*, 20 (4), 191-202.
- Singer, Peter (1991), "Research into Aging: Should It Be Guided by the Interests of Present Individuals, Future Individuals, or the Species?", in Frederic C. Ludwig (ed.), *Life Span Extension: Consequences and Open Questions* (New York: Springer).
- Slovic, Paul, Baruch Fischhoff, Sarah Lichtenstein, Bernard Corrigan, and Barbara Combs (1977), "Preference for Insuring against Probable Small Losses: Insurance Implications", *Journal of Risk and Insurance*, 44 (2), 237-58.
- Temkin, Larry (2011), "Is Living Longer Living Better?", in Julian Savulescu, Ruud ter Meulen, and Guy Kahane (eds.), *Enhancing Human Capacities* (Oxford: Wiley-Blackwell).
- Trotter, Griffin (2004), "Why Bioethics Is Ill Equipped to Contribute to the Debate About Prolonging Lifespans", *HEC Forum*, 16 (3), 197-213.
- Velleman, J. David (1992), "Against the Right to Die", *The Journal of Medical Philosophy* 17 (6), 665-81.
- Walker, Mark (2007), "Superlongevity and Utilitarianism", *Australasian Journal of Philosophy*, 85 (4), 581-95.
- Weinstein, Neil D. (1980), "Unrealistic Optimism About Future Life Events", *Journal of Personality and Social Psychology*, 39 (5), 806-20.
- Williams, Bernard (1973), *Problems of the Self: Philosophical Papers 1956-1973* (Cambridge: Cambridge University Press).