SCIENCE ETHICS AND SOCIETY

# Bioscience and the Good Life

Iain Brassington

BLOOMSBURY

Bioscience and the Good Life

# Science Ethics and Society

GENERAL EDITORS: **Professors John Sulston and John Harris**
(respectively Chair and Director of the Institute for Science,
Ethics and Innovation (iSEI) at the University of Manchester)

In conjunction with the Institute for Science, Ethics and Innovation, this series examines the major issues surrounding the impact of science and technology and the ethical issues generated by new discoveries.

Books in the series highlight the interplay between science and society; new technological and scientific discoveries; how they impact on our understanding of ourselves and our place in society; and the responsibility of science to the wider world.

Themes covered in the series include:

- Global justice
- The scope, limits and future of humanity ;
- Public health
- Technological governance
- Intellectual property
- Chronic poverty
- Climate change
- Environment

The series will appeal to policy-makers, academic readers and students in diverse disciplines including sociology, international relations, law, bioethics, physical and life sciences, and medicine.

**Books already published in the series:**

*International Governance of Biotechnology*
*Needs, Problems and Potential*
Catherine Rhodes
ISBN 9781849660655 (Hardback)
ISBN 9781849660778 (Ebook)

*Scientific Freedom*
Edited by Simona Giordano, John Coggon and Marco Cappato
ISBN 9781849668996 (Hardback)
ISBN 9781849669016 (Ebook)

*Beyond the Nuclear Family*
Daniela Cutas and Sarah Chan
9781780930107 (Hardback)
9781780930121 (Ebook)
Publication date: November 2012

*Global Health and International Community*
Edited by John Coggon and Swati Gola
ISBN 9781780933979 (Hardback)

# Bioscience and the Good Life

Iain Brassington

BLOOMSBURY
LONDON • NEW DELHI • NEW YORK • SYDNEY

**Bloomsbury Academic**
An imprint of Bloomsbury Publishing Plc

50 Bedford Square
London
WC1B 3DP
UK

1385 Broadway
New York
NY 10018
USA

**www.bloomsbury.com**

**Bloomsbury is a registered trade mark of Bloomsbury Publishing Plc**

First published 2013





**British Library Cataloguing-in-Publication Data**
A catalogue record for this book is available from the British Library.

ISBN: HB: 978-1-8496-6338-0
ePub: 978-1-8496-6339-7
ePDF: 978-1-7809-3093-0

**Library of Congress Cataloging-in-Publication Data**
A catalog record for this book is available from the Library of Congress.

Typeset by Deanta Global Publishing Services, Chennai, India

*Man, they said, is a very noble organism. We have dealt with other organisms so as to enhance in each its noblest attributes.*

*It is time to do the same with man.*

– Olaf Stapledon: *Last and First Men*

# Contents

# Acknowledgements

Parts of this book reproduce or adapt work that has been published or presented elsewhere. Chapter 5 reproduces material published as 'Body Art and Medical Need' in the *Journal of Medical Ethics* (32(1): 13–16, 2006); Chapters 8 and 9 recycle and adapt material published in *Bioethics* as 'John Harris' Argument for a Duty to Research' (21(3): 160–8, 2007) and 'Defending the Duty to Research?' (25(1): 21–6, 2011). The argument of Chapter 4 was presented as a paper at the 'Transforming the Human' conference held at Dublin City University in October 2011; the basic argument of Chapter 8 got its first outing at a seminar at the University of Bristol in December 2010.

All the material here has benefited greatly from discussion with colleagues in Manchester, especially in the course of CSEP and iSEI seminars run within the School of Law, and elsewhere. For fear of missing anyone out, I won't attempt to name names – but I'm very grateful to all of them for their comments and arguments and companionable sessions in the pub afterwards. I would also like to acknowledge the stimulus and support of the iSEI Wellcome Programme "The Human Body: Its Scope, Limits and Future."

# 1

# The Good of Bioscience

This is a book about ethics, and about the biosciences, and about some of the things that the former can tell us about the latter, and – in one chapter, at any rate – whether the latter can help us satisfy the demands of the former.

Ethics is conventionally thought to be about discerning right from wrong, and perhaps advocating the former. Bioethics is a branch of ethics that is about discerning right from wrong and perhaps advocating the former in relation to the biosciences and medicine. This is not surprising: most people's contact with ethics or ethicists comes, if it comes at all, when those ethicists appear in the media to talk about the rightness or wrongness of this or that action, or about how such-and-such an innovation should be regulated. Nor is the view entirely inaccurate: thinking (and sometimes talking) about the rightness or wrongness or the need to regulate things is an important part of the ethicist's job; and if ethicists are reluctant to take on the role of advocate, this is only because ethics as a field of study in the Western tradition is still only a couple of thousand years old, and debates about what is right plainly can't have been settled definitively in that short a time.

But it is not the whole story. While its true that ethicists do – at least sometimes – opine on rightness and wrongness, that isn't all there is to the discipline. (Elizabeth Anscombe famously argued in 1958, of course, that the view that ethics is concerned with the study and application of rules about rightness and wrongness is something of an add-on to what ethics was at the start of the Western tradition.[1])

Aristotle, for example, thinks that the scope of ethics is far wider than rule-setting. In the first couple of sections of the *Nicomachean Ethics*, he puts forward the hypothesis that all human activity is directed towards some good – the 'good for man' (I'll take the liberty of referring to it as the 'human good' from now on); and it follows from that that knowledge of that good is of great importance. The study of the human good, he continues, belongs to 'the science of politics'; political science 'lays down what we should do and from what we should refrain' in its pursuit. The investigation he undertakes within the *Ethics* is meant as a contribution to precisely that kind of political science.[2] The *Ethics* – and, implicitly, ethics itself – is concerned with securing the human good and (to take another of Aristotle's turns of phrase) *eudaimonia*: the flourishing 'good life'.[3]

We can take from this the idea that ethics may involve arguments about right and wrong; but that is as a part of a much less moralizing project, which is to get to grips with the contours of the human good, and make recommendations about what is and is not advisable in the pursuit of that good. It's this idea of ethics that's going to inform my claims over the coming chapters.

And because the biosciences – medicine, surgery, pharmacy, and so on – seem to take as their aim the human good as well, it would follow that ethics might have something to say about them. More, this something could well be richer than a set of claims about this or that procedure being permissible or impermissible: it would ask whether a given application of the biosciences really is conducive to the human good, and whether it's worth pursuing if the human good is its stated aim.

On the face of it, there is little question that the biosciences have contributed significantly to the human good. All of us could quite straightforwardly count ourselves as lucky to be alive today: at no previous point in human history have we had such power to combat illness and premature mortality. Illnesses that might have killed thousands every year not so many generations ago may well be almost trivially easy to treat now – assuming, that is, that they are still threats at all: smallpox is no longer a threat, and it's only politics that prevents the eradication of polio. Certainly, there are illnesses that we still struggle to treat: but there are reasonable grounds for optimism that at least some of them will become curable within the tolerably near future, and those that remain stubbornly incurable will at least be manageable. We understand a great deal about the causes and basis of a lot of illness; we can, at the very least, begin to map out what a definitive cure would have to do. For example, the Human Fertilisation and Embryology Authority recently (late autumn 2012) carried out consultation on 'mitochondrial transfer' – techniques for avoiding mitochondrial disease by, in effect, providing mitochondrial DNA transplants. Or – and this is my favourite current example of a future medical application for embryonic bioscience – we could use synthetic biology to realize the possibility of designing viruses to our own genetic blueprint that would destroy tumour cells without harming non-cancerous cells. (The AAV2 virus has been discovered to attack a large number of cancer cell lines without attacking healthy tissue;[4] it (or similar organisms[5]) could conceivably be engineered to a higher degree of specificity and efficiency.)

Nor is the potential capacity of the biosciences limited to reducing morbidity. As we'll see in later chapters, we could imagine them being put to use in other ways – for example as strategies to increase athletic or cognitive ability. Were it to be found that there is a correlation between, say, a version of a given gene and increased mathematical ability, we might imagine parents-to-be selecting embryos that carry that gene – or, slightly more radically, having it artificially inserted into embryos. Such procedures are science-fiction at the moment – but they needn't be forever. If we think that the human good would be served by making use of such technologies, we might well think that they would be worth

our while investigating. Such investigations might, some argue, even turn out to be obligatory.

There are some bioethicists who do not share such optimistic visions of what can be provided to us by the biosciences – and who are, in fact, positively suspicious of the promises made by them or on their behalf. I'll give voice to some of these arguments in the next chapter – though I will, in the end, dismiss them. I shall take it as more or less axiomatic that just about all possible future uses of the biosciences to serve the human good are morally permissible. The main possible exception to this rule of thumb is presented by 'moral enhancement', which I discuss in Chapter 7 and which I believe risks sacrificing an important aspect of human dignity[6] on the altar of too fragile a conception of goodness.

Still – if ethics is not just a matter of working out whether things are permissible or impermissible, assuming that the majority is permissible will not exhaust the scope of ethical debate. What remains to be seen is which applications of the biosciences stand a realistic chance of making a significant difference for the good, and whether that difference really does require bioscientific intervention. It's by asking this kind of question that I hope to be able to contribute to the bioethical debate that has come to the fore in the past few years concerning the promise that human lives can be made significantly better – and perhaps significantly different – by the application of recent advances in the biosciences and medicine.

My question will concern the ability of a range of possible things that might get done in the realm of the biosciences to contribute to 'the human good', or 'the good life'. I want to ask not so much whether this or that *ought* to be done, as whether there is much of a reason to do it.

And, though I reject the worries of the so-called 'bioconservatives' concerning the application of the biosciences, I shall also be distancing myself from some of the more optimistic voices. I shall be defending the claim that there is no particularly strong reason for any agent concerned with living the good life to pursue a significant number of the enhancements that are offered by the biotechnology of recent years or the expected near future; one doesn't really have to enhance to live the good life. In fact, more strongly, I suspect that there is at least one area in which we might well stand a better chance of living the good life by forgetting about what science might be able to do for us altogether: in this area, the benefits are not only illusory, but counterproductive. But that will have to wait: for the time being, it behoves me to say something more about how I understand the Aristotelian 'human good'.

## Understanding the good life

Given my aim of assessing the contribution that the biosciences might make to flourishing or the good life in future, it is tempting to try to give a full account of what the flourishing or good life would look like – not least because of the

suspicion that it would be difficult or impossible to make any particularly claims or assessments of the place of the biosciences in securing it without one. This is a temptation that I want to resist – not least because such a task would probably turn out to be huge, and to have a great many pitfalls. And I don't think it's necessary to provide a full account – one can make some general claims about at least some of the characteristics that a good life would display without having a complete account to hand; such a general account is plenty to be able at least to begin to say something useful about the biosciences.

What matters is that words like 'flourishing' have a meaningful place in the language-game that we play. To this extent, 'flourishing' need be no more substantial than a Wittgenstinian 'beetle in a box'.[7] Wittgenstein's metaphor for the relationship of words to ideas requires that we imagine a number of people, each of whom has a box into which none of the others can see; each agrees that the word 'beetle' refers to whatever is in the box. In this thought-experiment, it does not matter whether each box-holder actually does have the same thing as the others; it doesn't even matter whether there's anything in the box at all. However, we could imagine people comparing notes about their beetle's characteristics, and reaching a consensus about the appropriate use of the word. That is – they could work backwards from the role that the word plays in the language, and attempt to home in on giving it a precise extension.

It might be possible to say that certain things *aren't* a beetle quite easily. Indeed, it's only by excluding things that the word could be said to mean much at all, on the grounds that a word that means everything by that token means nothing. Some exclusions can be empirically informed (if we agree that beetles are characteristically found in matchboxes, they cannot be 75 feet tall); some conceptually informed (they cannot be vertebrate insects, given agreed usage of the words 'vertebrate' and 'insect'). Beyond this sort of thing, though, there can be debates about what makes a beetle that might never be definitively settled: the meaning of the word might be indefinitely contestable and contested – but this does not mean that the word is meaningless. It does carry some weight in language, and it doesn't seem to be completely incoherent.

As with 'beetle', so with 'flourishing'. The word's carrying weight in the language means we can start to look for something essential for it to mean if we want – but we don't have to; and we don't have to have access to its essence in order to be able to make claims about what it isn't, and to venture hypotheses about what is unlikely to contribute much to it. Questions of what constitutes flourishing (and what contributes to it) will then be matters for deliberation – which implies public debate and judgement according to standards of reasonability – and evaluation, based on how much from a list of desirable attributes and capabilities a life displays. That is to say, people can still articulate claims and hypotheses about the characteristics that one would expect within the good life; and these would still be scrutable by others. (This is an idea articulated and influenced by Martha Nussbaum's work.[8])

The public aspect of deliberation is, actually, of crucial importance. This is for a couple of reasons. The first is that, since there is no reason to believe that individuals are infallible in their assessments either of what flourishing is or of how best to achieve it, anyone who cares about maximizing his own flourishing will have a reason to accept the possibility that his understanding of what the word indicates is flawed in some way. And this means that anyone who takes his pursuit of flourishing seriously would have a reason to subject his particular understanding of the concept to refinement from other minds. He may simply have overshot the mark on some occasions, believing that something is necessary for a good life that public deliberation might lead him to discover is not. (One of the things that I'll be arguing over the first half of this book is that a number of putative enhancements are all very well, but that they represent, at best, an inessential way of living the good life.) The second is that there may be lapses and lacunae in his reasoning; he may have missed some aspect of the good life or how to live it. Again, the process of public deliberation may make these gaps clear to him and help him fill them.

That someone may overshoot the mark in respect of his expectations concerning the good life is important, because – for reasons that I'll elaborate in the coming sections – one of the things that may hinder flourishing is the inability to realize the ideas about himself that an agent has. But even without pre-empting that discussion, it is fair to say that there could be some expectations that a person has about a good life that are simply unreasonable, and that public scrutiny and deliberation would allow them to be 'tamed'. Some things, for example, might be straightforwardly impossible for anyone to realize; to include these things as criteria for a good or flourishing life is therefore foolhardy.

At this point, it might be worth forcing a distinction between things that are 'radically' impossible, and things that are not possible in a more workaday sense. I'll take 'radical impossibility' to refer to things that are beyond the scope of human possibility, and the latter to refer to things that are merely not possible given the resources currently available. For example, one of the things that might make a person's life better and bring her closer to flourishing is that she should be fairly fortune-favoured. Immunity from bad luck might be something that would contribute significantly towards a good life; and yet to suppose that there could be some kind of prophylaxis against bad luck stretches credulity. Immunity from bad luck is, therefore, radically impossible. Immunity from cancer, on the other hand, is merely not currently possible. At least at first glance, it would be foolish for someone to think that immunity from misfortune is a criterion for flourishing, even though it might well be desirable; this is the sort of thing that public deliberation would be able to regulate. I'll come back to this point in a little while.

Another feature of the claims and hypotheses about the good life that might be advanced is that they are likely to draw on a web of other claims about

human nature – each of which is, in turn, open to scrutiny itself. Aristotle provides a nice example of how this kind of argument might work at the start of the *Politics*: granted the hypothesis humans are, by nature, social animals, it follows (he reasons) that they will flourish within a social context, but not without one. Famously, any creature that has by its nature no social ties is either a beast or a god;[9] implicitly, anyone concerned to promote or pursue flourishing will have a reason to seek it within a social context. The same kind of reasoning could, *mutatis mutandis*, be produced in respect of any attribute whatsoever. Being animals, humans are by nature oxygen-breathing, which means that a reasonably good life will involve some, but not too much, oxygen; if humans were by nature aquatic, then the good for humans would involve water, and so on. Other claims about human nature might be more contestable, though the general point stands.

The point at the moment is simply this: to establish that it is possible to talk meaningfully about flourishing without having to have a complete positive account of flourishing present at hand. We can hope to be able to home in on it by means of deliberation and argument about what is reasonable – at the very least, by a process of elimination. As a placeholder, I offer here a hypothesis about what would constitute a good, or flourishing, life: the good life (for humans) is the life that could reasonably be held to be desirable (given a set of background claims about human characteristics). In slightly more detail, a good enough life is the kind of life that people could reasonably be expected to treat as acceptable; lives would be better the more they improve on being merely good enough, to the point at which the best life would be reasonably expected to be unimprovable.[10] With this – admittedly minimal – hypothesis as a guide, we can begin to try to work out what is and is not likely to be a feature of the good life: it is better than *good enough*, but need not be unimprovable. We can dismiss immunity to misfortune as being a necessary component of flourishing, on the grounds that since it is radically impossible, it is absent from any life – but since almost all lives are at least good enough, it must not, *modus tollens*, be required.

## Happiness and flourishing

Forced to give an initial hypothesis about a major characteristic of a flourishing life, one obvious response would be that it would be happy. This is quite possibly true: but we cannot deduce from that that increasing happiness is a straightforward route to a better life – and we can certainly dispute the idea that the best life is the happiest, in the sense of being the most pleasurable, life.

I think that it's fairly easy to establish this. Presumably, if flourishing were to be interpreted in terms of happiness, more happiness would mean more

flourishing. A person in a state of constant euphoria would be flourishing the most. So: imagine that someone is offered the opportunity to be wired up to some contraption that will stimulate the parts of his brain that produce happiness, and will either disable those parts of the brain that produce unpleasant feelings or modify them so that they are reliably and comprehensively overridden by happiness. In other words, imagine that he has the opportunity to spend as long as he wants with a guarantee of happiness and only happiness. Would this amount to an opportunity to flourish? I think that it would not.

One of the reasons for this is that the life on which he is about to embark is not recognizably human, and so does not seem obviously to fit the brief of providing human goodness: at least at first glance, it is odd to countenance the idea that catatonia and flourishing really inhabit the same sort of moral territory. I'm not sure that this is the best possible reason, admittedly; it seems to assume a great deal about human nature, and to be a little question-begging concerning the nature of the good life.

A more powerful reason is much more workaday, though, and it's that someone who was made no less happy despite any and all the things that might happen to him would have no obvious reason to eat, move to the bathroom before excreting, or attend to his bedsores. His life might be blissful, but it would likely as not be short and squalid. Not, admittedly, that he would care: but if people deliberating about the content of a hypothetically good life rejected this model as being the kind of life they would choose for themselves, it would be hard to maintain that they, rather than he, had made a mistake. So even if there is a reason for people to want this kind of pleasurable existence – and I'll accept that there is – it appears to me that pleasure will not suffice to establish flourishing. A flourishing life is one in which we might seek to eliminate hardships that we experience – but it is likely all the same to be one in which we *do* experience them. Those who would be lotus-eaters have made a mistake about what constitutes the good life.

At the same time, I don't think that it would be unreasonable for a person to think that a flourishing life would be one that doesn't attract the pity or contempt of others: but lying in one's own filth with a fixed grin on one's face could reasonably be hypothesized to do that. It's quite conceivably something that a reasonable person would want to avoid for himself; and it's therefore not a part of the good life.

So: a good life is not reducible to being a pleasurable life. While happiness may be an important indicator of flourishing, it ought not to be taken as the sole or sufficient indicator of flourishing. Is it a necessary criterion, though? It makes perfect sense to talk about the rosebush in the garden as flourishing, and it takes no pleasure from it. I think it likely that pleasure would be a part of a good life, inasmuch as that a good life seems to be by definition one that is better and more desirable than neutral. A neutral life would be one the end or continuation of which was a matter of complete indifference to the one

living it. Since I'm reluctant to commit to the idea that pleasure is the only characteristic of a good life, it might turn out that the goodness that pleasure provides is matched by the goodness that something else provides; and so someone could think his life worth living despite the absence of pleasure. (Perhaps a grim sense of duty for its own sake would fit the bill; or it might be that someone could say to herself that she is neither happy nor miserable, but wise, and that wisdom is worth having.) For sure, a life like this may be *better* for having pleasure in it, but there is no real reason to suppose that its absence is a deal-breaker.

## The importance of projects

'But what do we do now, now that we are happy?'[11] asks Estragon in the second act of *Waiting for Godot*. It's not a bad question, really. Granted that a miserable life would not be a good life, it doesn't follow that happiness is enough; so what else has to be present to secure the good life? I contend that a necessary part of the good life is the presence of projects of one sort or another.

Humans are temporal creatures. Unlike at least some other animals, we have a sense of ourselves enduring, and the ability to survey our lives as lasting throughout time. This means that we have a stake in the future, and a corresponding interest in *our own* futures. Importantly, our ability to see our lives as stretching out into the future – more: our inability *not* to see that life stretching out into the future – means that we are susceptible to boredom. For as long as he could conceptualize himself as persisting for an indefinitely long time into the future, someone without any projects would not stand to derive any satisfaction from that life; rather, it would yawn ahead of him as something, at best, to be endured. And a life that does not promise any grounds for satisfaction is quite plausibly less good than one that does. (Being alive is not its own reward.) For this reason, it would seem that a person who lacks projects would not be one to whom we could ascribe flourishing; rather, his life would seem to be strangely listless.

(If this argument holds water, it implies that it is better to have things that one would like to do but which are yet undone, than it is to have realized all of your projects and have nothing more that remains to be done. I accept this implication, and it'll prove to be an important consideration when it comes to analysing the good of significant life extension when we look at that in Chapter 4.)

Relatedly, someone who has no projects at all, however minimal, would seem to have no particular stake in the future. Lacking any such stake would mean that he would have no obvious reason to be anything but indifferent about surviving to see that future.[12] Granted, he would be relieved of certain burdens – notably, the worry that his projects might founder. Yet this is only a

great burden on the assumption that the good derived from a project depends on its completion. Such an assumption is dubious: though a hike 'contains' the idea of its completion, the hillwalker is interested in the walk rather than its termination, for example; the concertgoer wants to hear the music being played. It is not therefore obvious that seeking protection against projects being unrealized by having as few as possible is really the best move; a life lived according to this rubric would likely be more indifferent than particularly *good*.

Recognizing a role for projects helps us understand the place of happiness in the good life, too. The thought here takes its lead from a suggestion made by Bernard Williams in his 'integrity objection' to utilitarianism. Williams is dissatisfied with utilitarianism's claim to be concerned with the promotion of happiness as the end to which actions are and should be directed, claiming that to emphasize happiness as an end misses something important. 'One has to believe in, or at least want, or quite minimally, be content with, other things, for there to be anywhere that happiness can come from',[13] he reasons. We may derive happiness from the pursuit and occasional realization of projects – but that is not the reason why we embark upon a project or live that kind of life to begin with. Rather, the happiness generated comes from living a certain kind of life. Though the pursuit of projects is a source of happiness, neither hillwalkers nor concertgoers do what they do *for the sake of* being made happy by it; if they did, they ought to be indifferent between that activity and a neurological implant that will assuredly beam just as much happiness directly into the brain. I've already argued that that wouldn't be a reasonable thing to want; but we can now add to the reasons I produced earlier another: that such an implant would strip the euphoric life of its recognizable 'mineness'. Happiness is undeniably a good thing – but it is found in a certain life. The point of being a hillwalker is the walking of hills; the pleasure derived is not (and is not supposed to be) substitutable for other pleasures.

So we can begin to understand Estragon's question. Becoming happy and then wondering what to do shows that happiness up as the sham it is. But while there's nothing to be done (he and Vladimir tell us as much twice each during Act 1), that's all there is; and his life won't be a good one. On the other hand, if there is something to be done, and that something is something with which Estragon can identify to the extent of being able to say that it's what his life, or this part of his life, is *about*, then he stands a chance of being happy 'by accident' – and his life stands a chance of being good.

## Function and the good life

What else may we come to expect to see in a good life? It would be a surprise not to see some recognition of physiological function there, inasmuch as that

it's tempting to think that a good life will be one in which the body behaves as bodies ought – but even here, things aren't necessarily as straightforward as we might think at first.

The problem is that talking about bodies functioning more or less well presupposes a great deal about the standard that we're supposed to be using to measure that function. Certain things might be fairly easy to agree: a wound that won't heal is pretty straightforwardly a sign of dysfunction, for example. But other things are harder to agree. Too naïve an account of function and dysfunction in relation to flourishing might well struggle to make sense of the fact that a lot of people seem to flourish despite some ostensible physical dysfunction, and would not obviously benefit from bioscientific intervention. For example: imagine that a person – call her Joanna – is deaf. We can stipulate that this is because of some congenital characteristic, and that modern bioscientific insights have given us the ability to give her the ability to hear. But it is not a given that Joanna's deafness prevents her from flourishing: there are plenty of deaf people who believe that their lives are different from, but no worse than, the lives of their hearing counterparts, for example. Joanna might believe that being able to hear would be a boon; but she might not. Either way, it is not obvious that she would have to think it worth the effort, and she might even think it undesirable. What will make a difference to whether or not any such intervention is warranted will have a lot to do with whether she misses the capacity to hear; this will have a lot to do with the kind of life she wishes to lead; and much of that picture will, in turn, be informed by the kind of projects that she has for herself. If it happens that Joanna does not miss hearing, we can say that her deafness is in one sense a dysfunction, without having to admit that it makes any particular difference to her flourishing. Therefore flourishing can't depend closely on function in this tight, and slightly scientistic, sense.

Or imagine that Brian is infertile. On certain accounts, this is straight forwardly a dysfunction, and a characteristic that it would be better not to have. However, it is not clear why a bare description of this characteristic would allow us to say anything about how his life could be improved: there is no obvious logical link between a barely descriptive statement about a person's physical characteristics or attributes, and a statement to the effect that these characteristics or attributes matter. It is by means of appealing to certain projects – notably, the project of raising a family that is genetically related to him[14] – that we can draw a link between a bare description of function, and Brian's flourishing. If infertility impacts on flourishing, that is, it does so only in the context of a particular kind of parental project that is not realized. Without such a project to make sense of that desire, it might be wholly unclear why providing it would make any improvement to his flourishing.

Neither of these examples deals with illness; but it might be possible to extend the account even there, to help us work out what is meant by words

such as 'health' and 'illness'. Havi Carel is insightful here, suggesting that the healthy body is 'transparent', and that

> while digestion, fluid balance and muscular performance are going well, we do not experience them consciously. They silently and invisibly enable us to compose symphonies, have coffee with friends and argue about politics.
>
> It is only when something goes wrong with the body that we begin to notice it. Our attention is drawn to the malfunctioning body part and suddenly it becomes the focus of our attention, rather than the invisible background for our activities.[15]

The difference between illness or certain kinds of injury and a dysfunction such as infertility is that the former interfere with our lives on a much more pervasive level, and so their impact is felt in the context of much more mundane projects and activities that we are used to pursuing and performing 'at a certain speed, in a certain way' – things as mundane as chopping vegetables, or moving from one room to another. Carel suggests that

> [i]llness distances us from the biological body, which becomes alienated and erratic, the source of pain and disability. The lived experience of this body becomes painful, unstable, treacherous. The distance from the biological body is not normally available to us while we are healthy. Illness (as well as other kinds of physical alteration) removes the body's transparency and problematizes it.[16]

It's in doing things that illness often makes its presence felt. As such, there's a connection between illness and projects – or, more widely, doing things – just as there is between dysfunction and projects.

The naïve functionalist account would also potentially struggle with the fact that a lot of other people may apparently be helped to flourish by some intervention notwithstanding that that particular intervention does not address any particular dysfunction. If this help is not illusory, it undermines the possibility of drawing too neat an association between function and flourishing. Instead, we need to be able to give any such association a context; appeals to projects can do this, and also give us a reason to think that there are means in which the biosciences can help the non-dysfunctional to attain higher levels of flourishing. For example, in Chapter 5, I'll consider situations in which a people might choose to take advantage of biomedical interventions to make their lives better without there being anything much wrong with them to begin with.

Of course, projects might not be the whole story. People do have a concept of themselves, and they might well be able to point to a sense of personal integrity as being important in their assessments of their flourishing. For example, women who have had a mastectomy are frequently offered reconstructive surgery, and one of the rationales offered for this is that it is of psychological benefit – that it may not assist with the cure for cancer, but that it does help with healing.[17] A procedure such as gender reassignment surgery could also draw on this kind of

argument: if a person is convinced that they do not 'fit' their sexual phenotype, this can cause a severe diminution of flourishing – a diminution that we might legitimately hope to halt or reverse; and no further project would be necessary to make sense of how such an intervention would contribute to flourishing. So we might be able to enhance a person's level of flourishing by modifying their body in some way in accordance not with any particular project, but rather in accordance with their vision of the kind of person that they are on a more existential level.

Still: the relationship between physical considerations such as (but not limited to) health and flourishing is potentially quite complicated. A person might be flourishing notwithstanding some characteristic that could, from at least some perspectives, be seen as dysfunctional; and their flourishing could be improved despite the absence of any dysfunction.

This appears to count against placing too much stress on the distinction between therapy and enhancement. If the justification of the use of the biosciences is to improve flourishing – to make lives better – then it ought not to concern us too much whether the intervention in question is aimed at rectifying some dysfunction or illness, or improving an already functional state. There are projects that a person might have that are hindered by the bare realities of their body even when that body is not dysfunctional in any relevant sense; and if having one's projects hindered is capable of denting a person's flourishing, and if the realization of those projects can be assisted by making use of the products of the biosciences, then that is at least *prima facie* all we need for there to be a warrant for using them. It makes no difference whether we are restoring or repairing in a person who lacks the capacities that most people do have, or whether we are giving them capacities that most people don't. I will, however, return to the therapy/enhancement distinction in a little while.

## The reasonable expectation standard

I've already indicated that someone without projects, though probably not living anything like a good life, would at least be able to ensure that she is never disappointed by their not having come to fruition. Similarly, a person with no particular conception of themselves as having this or that kind of life cannot be let down by its not having been realized. And it does seem reasonable to suppose that a life that is marked by disappointment will be less good than one that is not so marked. As such, disappointment is corrosive of flourishing. It is, in fact, 'positively' undesirable: though the value of a project and its contribution to the good life doesn't depend on its having been realized, still, to be positively disappointed might well mean that one flourishes *less*. So this means that there is at least a possibility that a life without projects, though barely worth the

candle in its own terms, might be just-about better than a life characterized by systematic disappointment. Still, this is hardly a ringing endorsement of such a pessimistic life. More importantly, a person who withdrew from life in order to guarantee not being disappointed could, I think, be said to care less about the good life than about securing the least-bad life; as such, he'd be satisfied with a life that's just about good enough – not as good as some, though not as bad as others. I think we can assume that a person could justifiably aim higher than that.

I have already claimed that the nature of the good life is a matter for public debate; and, this being the case, it is likely also that the means by which we might try to attain it is also quite properly material for public debate; and just because a given project or view of the good life is pursued sincerely and in earnest, it doesn't follow that it can't be misguided. To this extent, claims about the good life can be treated as hypotheses, at least in principle falsifiable. For instance, in the example I gave above, Brian attaches a great deal of weight to becoming a parent and understands this to require that the child be formed from a gamete that came from his own body. But it's possible that he's mistaken in identifying this as a requirement: it's possible that he could be a parent in the most meaningful sense by means of adoption. His beliefs about the impact of fertility on the goodness of his life could have the status not of axioms, but of hypotheses, capable of being abandoned.

The lenses through which claims about the good life or the means to achieve it can be examined are several and various, but all contribute to an assessment of whether the proposed criteria for a good life are reasonable. So, for example, we could counsel against the simultaneous pursuit of two mutually incompatible or antagonistic projects (albeit projects that are in their own right perfectly admirable) on the grounds that it is impossible to do both, and so not reasonable to try. There could be times, then, when a person's notions of the good life and how to achieve it are tangled; it might therefore be necessary to prune the thicket. Performing this task might be easier with the help of others. If we're trying to secure flourishing, it might make sense to rein ourselves in from time to time; and it might make sense to allow others to help, too. Granted that we know ourselves to be fallible, we have at least a *prima facie* reason to allow our visions of the good life to be honed by public scrutiny.

If we take flourishing seriously, there is no particular reason to be chary about correcting (or having corrected) false claims about what constitutes a good life, or misguided or straightforwardly loopy strategies that people might adopt to live as good a life as possible. It would be appropriate – or, at the very least, not inappropriate – to recommend that some ideas of the good life and the projects undertaken to achieve it should be unceremoniously ditched. (There's an Aristotelian tinge to this claim, too: in the *Ethics*, Aristotle specifies that ends – such as the good life – are not subjects for deliberation, but the

means to those ends are: '[w]hat we deliberate about is practical measures that lie in our power'[18]; and 'while the end is an object of wish, the means to it are objects of deliberation and choice'.[19])

As a general rule, a hypothesis about what makes for a good life is reasonable if it is not self-contradictory, or false, and possible. Hence thinking that a good life would have to be free of bad luck would invite falsification by having it pointed out that plenty of people have claimed, without obvious error, to be living a good life notwithstanding the occasional misfortune. To be free of illness is a more plausible demand for the good life, but, again, there's no reason to disbelieve those who think that their lives have been (or at least could have been) good despite illness,[20] and freedom from illness would probably have to be recognized as setting the bar too high given that humans are, by nature, lumps of biology in a biological world.

The potential pitfall of this strategy is that it rebounds too far, inasmuch as that we might be tempted to ask whether 'reasonable' has to mean the same thing for all people. What is reasonable for one person to expect from life may not be same as what is reasonable for someone else to expect because of prevailing contingencies.

For example, a person living in a refugee camp somewhere after having fled war and famine might have all kinds of desires and expectations about what the future could bring her. Some of these might be more fantastical than others, but the more pressing problem may well be that even the less fantastical ones are very difficult to realize, and so might count as practically impossible for someone in her situation – and so it'd appear that we'd have to dismiss her conception of the good life as unreasonable. Conversely, what it *is* reasonable for her to expect and desire given the practicalities of the matter may still only generate a life that would strike many in the developed world as being, at best, minimally tolerable. This would seem to generate the highly counterintuitive conclusion that we ought to abandon a picture of a good life for her that people in wealthier parts of the world would take for granted.

However, this objection is surmountable. Assuming that the life our refugee lives is the only one she knows, *she* might think that it is reasonable to mould her expectations to fit it – but this is not something that everyone has to accept. I've already claimed that a person might be mistaken in respect of what they think they could reasonably expect from life, and others could – presumably – assert that such a life falls below a standard that is acceptable as reasonable, and that she is in actual fact entitled to a set of expectations from life that looks significantly different.

An important consideration here is that there's a plausible difference between reining in one's visions based on some understanding of what it is to be human, and reining them in based on some understanding of the contingent conditions of the rest of the world. Insisting on pervasive good fortune or invulnerability to illness as a criterion of the good life is unreasonable, given background

claims concerning humans' place in the world; insisting on security and access to a certain level of medical care in the event of illness, though, isn't – and we know that it isn't unreasonable, because it's something that is trivially possible for a large portion of humanity already. So while someone in the poorest parts of the world might have a different idea about the details of the good life, the essentials – being well-fed, healthy, sheltered and so on – are likely to be fairly constant.

This is as much as to say that, even if someone's imagination about what constitutes a good life is, to our eyes, limited by experience, it doesn't follow that we have to accept those limits. After all, Mill's pig's failure to appreciate Socrates' life doesn't mean that Socrates' isn't better.[21] Even if someone thinks that their life is roughly as good as it could reasonably be, they might be mistaken.

I'm even drawn to the idea that a few moderately unreasonable desires about the good life might be worth retaining, just as long as they are recognized as not being foundational to a good or good enough life. Partly this is because the pursuit of things that are currently practically impossible might give some people the kind of project that I've already suggested would play a big part in anything like a good life. But it's also because many of the good things we enjoy now came about as a result of people having slightly unrealistic desires, and working to realize them; and there is every reason to suppose that a great number of future improvements will also be generated because of people's unwillingness to be reasonable.

In fact, there is no reason to suppose that what counts as required for a flourishing or a good life is a static target. By this, I mean that people living in previous ages would presumably have had occasion to think their lives good – and yet were we offered the chance to abandon our lives in favour of theirs, we would quite possibly find that they did not meet our intuitions about what was required for a good life: a life without the benefits of running water, sunglasses, effective antibiotics, and so on might be bearable, but it would be significantly less good than the lives we have now. This might mean that some of the things we take to be necessary for a good life are actually not; but it's also possible that, as new technologies become available and become familiar to us, they can alter our view of what a human life is and can be, and of what would be good therein. Future generations may look back on us and wonder how we managed to get by without some piece of technology that they take for granted.[22]

There is plenty of scope for a 'race to the top' in which, as time goes on, what was accepted as a reasonably high standard of life at time $t_1$ comes to be accepted as being only barely acceptable at $t_2$, and what was considered to be a rare treat comes to be seen as more of a necessity. One might care to analogize this kind of progress in terms of plumbing: not so many generations ago, a house built without indoor plumbing would have been seen as perfectly acceptable; it would not be now – which is as much as to say that today's

judgements about what counts as minimally constitutive of a good house are far in excess of yesterday's (and tomorrow's may be different again). In each case – plumbing or anything else – what was once extraordinary will become, with time, expected and attain the status of a necessity; what once would never have been classed as a central part of the good life may be today. Adam Smith seems to be making a similar point when discussing the difference between luxuries and necessities in *The Wealth of Nations*:

> By necessaries I understand, not only the commodities which are indispensably necessary for the support of life, *but whatever the custom of the country renders it indecent for creditable people, even of the lowest order, to be without*. . . . Under necessaries, therefore, I comprehend, not only those things which nature, but those things *which the established rules of decency have rendered necessary to the lowest rank of people*. All other things I call luxuries, without meaning, by this appellation, to throw the smallest degree of reproach upon the temperate use of them.[23]

Some things might be necessary based on an appeal to physiological need; but necessity is broader and richer than that: it can be shaped by cultural norms. Likewise, what's good – as opposed to what is just about good enough – is an expansive set of things.

(Martha Nussbaum has a telling anecdote in this respect:

> A recent comparative survey of Indian widows and widowers established that the widows almost always described their health status as "good" or "excellent," even though, from the medical point of view, they were doing rather badly by comparison with the males (who tended to report numerous complaints and to rank their status rather low). Ten years later a similar survey found no significant change in the women's actual health status. But their perceptions of health deficiencies had improved, as a result of education and public information. Now they perceived more nearly what the doctor did, and understood that there was an available norm *vis-à-vis* which their own condition was seriously deficient.[24]

The telling aspect here is that the widows came to see that there was an available norm that they had not seen before. Put another way, their view of what it is reasonable to accept seems to have shifted. Daniel Everett hints at something similar in his talk of the attitudes of the Pirahã Indians of the Amazon to death:[25] it's not that they see dying in one's thirties as anything other than bad; but in a non-technological culture such as theirs, it's also an inevitability, and so does not represent a particularly gratuitous violation of any natural propriety. People that have no idea that others elsewhere routinely live at least twice as long as they will not have the same notions of what counts as the unacceptable attenuation of a life.)

A world in people can strive for more than is required by a minimally acceptable standard of flourishing – or even a standard of flourishing that is largely accepted as fulsome – is one in which at least some will be able to invest in innovation and supernormal benefits. If and as these benefits begin to

filter through to the rest of the population, they will gradually inform what is accepted as a minimum reasonable degree of flourishing, and thereby – as if by osmosis – drag the general level of humanity's flourishing with them.

(Before moving on, I'd like to return to the dispossessed refugee from a little while ago. I think it's reasonable to say that a part of the good life will be a sense of community with others. This claim is not based on any supposition about humans being social by nature; rather, it has to do with the much more straightforward idea that a life in which one feels alienated from the lives and projects of others is likely to be less good than one in which one does not. Alienation, therefore, is detrimental to the good life; community isn't. And so it might turn out that the lives of the poorest are made worse than they otherwise would be, and less likely to qualify as being good in any sense, by the mere knowledge that they are already significantly worse than those of others.

(As such, radical disparities in welfare, irrespective of how well-off the worst-off are, ought to concern any ethicist who is concerned not just to *understand*, but to *promote* something like flourishing. One response to this might be to say that it'd be better for the worst-off never to learn how well-off the best-off are – but this kind of response is not only morally questionable on the grounds that it's more than slightly patronizing; it is also perplexing in a non-moral sense. After all, being motivated by this sort of concern looks like being motivated by some sort of concern for the level of flourishing of the person about whom we are thinking. But if we have that concern, it's very strange that we should think it acceptable to promote contentment in terms of blissful ignorance rather than the kind of good life that we would accept for ourselves.)

## Rebooting the therapy/enhancement distinction

There are obviously some aspects of the good life that the biosciences cannot help us realize. But it is equally obvious that they could easily be enlisted to help us maximize our flourishing in respect of others. The emergence of a previously unknown illness may well represent a threat to the flourishing of those infected with it; it's thanks to the insights of the biosciences that we stand a reasonable chance of identifying the weaknesses of the virus that carries it, and minimizing its impact: SARS gives us a nice example of the way that a new and threatening illness can be identified, controlled and – all being well – eliminated within a very short time. Progress in the biosciences allows us to be much better at regenerating damaged tissue than nature is: a timely skin-graft may not reverse the effects of the burn entirely, but it can reasonably be expected to be far better at repairing the damage than would unassisted nature, for example – and as time goes on, the complete reversal of the injury may become possible. Ways to minimize or reverse the effects of Parkinson's disease,

multiple sclerosis, diabetes and any number of other conditions that limit the length and quality of lives are increasingly available.

But if we're prepared to endorse making use of the biosciences to *repair* people for the sake of a good life, why not take the small extra step and endorse using those same sciences to improve them from an 'undamaged' state? That is: why not leap from therapy to enhancement?

A very powerful argument in favour of enhancement is to capitalize on the conceptual proximity of the two. The essence of the comparison is captured by John Harris in *Enhancing Evolution*: '[m]ost of what passes for therapy', he says, 'is an enhancement for the individual relative to her state prior to therapy.'[26] This does seem correct; certainly if we take seriously his assertion that an enhancement is 'anything that makes a change, a difference for the better'[27] – which is at least *prima facie* about right – then it would seem to follow that therapy is indeed a kind of enhancement. (I shall have more to say about Harris' definition of enhancement in Chapter 3.) Since noone would deny that therapy is a good and desirable thing, the argument runs, and if there is no significant difference between therapy and enhancement, they are bound to agree that enhancement is also a good and desirable thing.

However, while this argument does appear to show that to make a therapeutic intervention is to make an enhancing intervention, it does not show that there is no distinction between therapy and enhancement. One can accept that therapy is enhancement without having to accept that there is no distinction between them, just as one can accept that a square is a rectangle without having to accept that there is no distinction between the two: there plainly is. The relationship between a square and a rectangle is that the former is a species of the latter: in thinking about a square, one is inescapably thinking about a kind of rectangle – to wit, an equilateral one – but one could think about rectangles that are not squares. In the same way, in thinking about therapy, one is thinking about a species of enhancement; but one could think about kinds of enhancement that are non-therapeutic nonetheless. Since one can coherently think or talk meaningfully about enhancements without at the same time thinking about therapy, it follows that there must be at the very least a logical distinction between the two. And since there is a logical distinction between the two, even if they turn out to have a great many of their ethical characteristics (or even all of them) in common, the arguments produced in respect of one will not automatically translate to the other.

So if there is a distinction, what might it be? On the basis of the argument I've just put, I think that it's possible to offer an answer to this. I've accepted as axiomatic the idea that 'the good life', or 'flourishing' is desirable, but that its precise constituents are unclear. This unclarity needn't be a problem, so long as we are prepared to accept the idea that a good life is simply one that does not consistently fall below a particular standard of reasonable acceptability. And this is all we need to reboot a therapy/enhancement distinction (henceforward,

'T/E distinction'). The word 'therapy' captures procedures that aim to improve an aspect of a life that would not ordinarily be considered to be minimally acceptable to the point at which it is.[28] Enhancement is whatever gives a person more than is minimally reasonably acceptable.

I think that one of the attractions of this account is that it tracks the way the words actually do operate in everyday life. When people talk about therapy, they generally mean something like the sort of enhancement that makes a person tolerably well. 'Enhancement' is reserved for the things that go beyond that.

Harris has a response prepared for this challenge:

> [T]hose who wished to cling to the therapy-enhancement distinction could simply say "yes, therapy and enhancement sometimes amount to the same thing, but not always." But now consider that many routine therapies – vaccination, for example – are enhancement technologies for the simple and sufficient reason that they enhance and the moral significance of the distinction as well as its utility collapses. It does not draw either a morally significant or an explanatorily significant distinction and so fails utterly to be useful.
>
> [. . . M]y point is that treatments or preventive measures which protect humans from things to which they are normally vulnerable or which prevent harm to that individual by operating on the organism, by affecting the way the organism functions, are necessarily also enhancements [. . . T]he distinction between therapy and enhancement, between protection and improvement, cannot be coherently or consistently maintained.[29]

I think that it is a bit soon to say that the significance and utility of the distinction has collapsed – for as long as the distinction is there, its significance and utility depends on what you want to do with it, and insignificance and disutility will not demonstrate non-existence. But, for the time being, it might be enough to point out that Harris is here expending a great deal of energy pushing at an open door, since noone has to deny that therapy is a kind of enhancement to maintain that there's a difference between them. What noone is committed to accepting, though, is that all instances of enhancement are at the same time instances of therapy – unless one subscribes to the bizarre idea that since lives can always be made better, they are always to that extent pathological. (In a similar vein, when Ruth Chadwick says that 'try[ing] to explain enhancement through making an opposition with therapy is far from a simple matter'[30], she is right – but she is also worrying about the difficulty of something that we don't have to do.) Further, there is something potentially problematic with Harris' stipulation that vaccination is therapeutic, just because it is possible to lead an acceptably good life without it (and it might be better to do so all else being equal: we no longer vaccinate against smallpox, and look forward to the day when we no longer have to think that polio vaccination is worthwhile). Vaccination reduces the requirement for future therapy, but it is not therapeutic itself; in the same sort of way, you don't rescue someone by teaching him to swim, although you certainly do reduce the chance that rescue will be required.

This makes a difference to claims about the contributions that can be made by the biosciences to the good life. Notably, there is likely to be a range of characteristics recognized as being necessary for a life agreed to be good enough, and another range of characteristics that, while capable of making lives better than good enough, is not necessary. Without characteristics of this latter sort, we could still have a life worth living. We would still have a reason to pursue enhancements – a superb life is, naturally, better than a good one – but we would nevertheless have an acceptably good life without them.

## Closing the distinction?

Admittedly, though, the fact that there's a logical distinction between therapy and enhancement isn't sufficient to tell us that there's much of a practical distinction, or that any distinction is significant; and a range of thought experiments can be produced that are directed at showing that whatever significance it has is, at most, minimal. Harris' argument purports to show that anyone who rejects therapy is *pro tanto* rejecting enhancement; but a much more powerful form of scepticism about the T/E distinction comes from a scenario that Nick Bostrom and Toby Ord have generated. This aims to show that a refusal to endorse a particular enhancement is effectively a refusal to endorse a particular kind of therapy. If this conclusion is correct, it does serve as evidence that the notional distinction can be narrowed or maybe even closed.

Bostrom and Ord's scenario is this:

> Disaster! A hazardous chemical has entered our water supply. Try as we might, there is no way to get the poison out of the system, and there is no alternative water source. The poison will cause mild brain damage and thus reduced cognitive functioning in the current population. Fortunately, however, scientists have just developed a safe and affordable form of somatic gene therapy which, if used, will permanently increase our intellectual powers just enough to offset the toxicity-induced brain damage. Surely we should take the enhancement to prevent a decrease in our cognitive functioning.
>
> Many years later it is found that the chemical is about to vanish from the water, allowing us to recover gradually from the brain damage. If we do nothing, we will become more intelligent, since our permanent cognitive enhancement will no longer be offset by continued poisoning. Ought we try to find some means of reducing our cognitive capacity to offset this change? Should we, for instance, deliberately pour poison into our water supply to preserve the brain damage or perhaps even undergo simple neurosurgery to keep our intelligence at the level of the status quo? Surely, it would be absurd to do so. Yet if we don't poison our water supply, the consequences will be equivalent to the consequences that would have resulted from performing cognitive enhancement in the case where the water supply hadn't been contaminated in the first place. Since it is good if no poison is added to the water supply in the present scenario, it is also good, in the scenario

> where the water was never poisoned, to replace that status quo with a state in which we are cognitively enhanced.[31]

Their case does look powerful: an increase in intellectual capacity is an increase in intellectual capacity, irrespective of whose intellect it is. If it's good to increase Smith's capacity to offset the poison in the water, it ought also to be good to the same extent to increase Jones' even though – in Jones' world – the poison isn't there and she faces no counterfactual diminution of her intellectual ability; after all, what we're talking about in both cases is providing something good and desirable. By that token, if it's absurd (at best) to resist giving the intelligence-preserving concoction to Smith, it looks as though we're going to be committed to saying that it's absurd (at best) to resist giving it to Jones. So it's hard not to get swept along with the scenario generated by Bostrom and Ord, and they do make a powerful case for the moral equivalence of something that we would intuit to be a therapeutic intervention in Smith's case and bare enhancement in Jones'. However, they make a couple of assumptions that are not bulletproof.

Presumably, if we think that we ought to use the concoction to counteract the poison, we have to think that it would be bad for humanity to lose this amount of cognitive capacity. But the badness of losing something is not the same as the badness of never having had it. Let's imagine that the poison, unmediated, would bring about a 25 per cent drop in Smith's cognitive capacity. It's reasonable to want to prevent or reverse this. The genetic therapy would achieve that. But in respect of Jones, it represents no prevention or reversal at all. We would be helping Smith to keep hold of something that he could reasonably be expected to value but would otherwise lose; but in Jones' case, we would simply be helping her to get more of something that she would have kept anyway. The difference is roughly that between helping Smith keep the burglars away so that he doesn't lose his stereo, and helping Jones get an upgrade to hers. Both end up better off than they might have done had it not been for our intervention; but accepting that we have a reason to help Smith does not commit us to accepting that we have the same reason to help Jones. In both the Bostrom/Ord water example, helping Smith would be an example of the kind of intervention directed at guaranteeing an aspect of life that is at least minimally reasonably acceptable – the kind that is, though *stricto sensu* enhancing, captured by the word 'therapy' in ordinary language. Helping Jones would be the kind of intervention not directed at guaranteeing any such kind of life – the kind that, in everyday English, is referred to as just plain old, unqualified, enhancement.

This sort of consideration tallies with the other assumption implicit in the Bostrom/Ord example, which is that the purpose of interventions is to secure improvements rather than to secure some standard of flourishing. If we assume that Smith was living the kind of life that we might take to be as flourishing

(or, minimally, flourishing in the relevant sorts of way) as could reasonably be expected before the problem with the water, adding the treatment in its wake would both provide him with a benefit and restore him to a reasonably good life. It might also be highly admirable for us to provide Jones with a benefit, too – but her pursuit of the good life is not threatened, and so our reasons to act do seem to be less urgent in her respect. Interventions that make a reasonably flourishing life better strike me as being properly described as bare enhancements; those that help ratchet a life that is not reasonably flourishing closer towards that standard strike me as being properly described as therapeutic.

## The structure of this book

There are two kinds of good life. What I have been considering so far is the eudaimonistic vision of the good life. On this view, to live a good life is to live one's own life well – to approach it as a sculptor would approach a piece of marble, in the hope of getting something pleasing out of it, with as few rough or ill-proportioned bits as possible. But there is another view of how to live a good life – a view that says that one lives a good life when one's actions reflect one's obligations to others.

For the most part in this book, I shall be considering the good life under the light of the first understanding. I'll take it as read that therapeutic interventions and products of the biosciences are contributory to the good life, and ought therefore to be endorsed largely without qualification. The one exception to this claim comes in Chapter 7, in respect of so-called 'moral enhancement' – the idea that we may or ought make future people better by making them better behaved. I shall claim, much more strongly, that there is a positive reason not to pursue such technologies that borders on making them morally wrong. Beyond that, I'll be attempting to assess how, and whether, a number of kinds of enhancement contribute to the good life. I am fairly sceptical about the ability of a number of things promised by the new biosciences to make our lives all that much better. On the other hand, there is a few things to which we could turn our biomedical and bioscientific skills that *do* make our lives better. Interestingly, I shall argue that the more whimsical and playful the application, the higher the chance enhancements have of improving our lives.

In Chapters 3, 8 and 9, I shall focus my attention away from what it is wise or advisable for us to pursue, to what it is obligatory to do in respect of the new biosciences. Here, I shall claim that there are very few obligations at all. It is perhaps admirable to make full use of the biosciences – and it is certainly rarely wrong – but I shall defend the position, *contra* some, that there is nothing like a duty to enhance, and that there may not even be a duty to carry out many – or perhaps any – kinds of medical research.

Before I get that far, though, I shall spend the next chapter giving a brief survey of the most famous arguments against the radical use of the new biosciences, and why they don't work. The interesting debate is not between those who think that enhancement is permissible and those who think that it is not; it is between those who think that enhancement is permissible and those who think that it is obligatory.

## Notes

**1** Anscombe, 1958, *passim*.
**2** Aristotle, 1976, pp. 63–4.
**3** Ibid., p. 73.
**4** Voigt and Keasling, 2005, p. 305.
**5** Ramachandra et al. 2001, *passim*.
**6** 'Dignity' is a word that I half-think should be placed on the index: it is a word that is at once much too loaded, much too under-defined, much too malleable, and much too much else besides. On this, I'm wholly in agreement with Alasdair Cochrane. Nevertheless, it does occasionally have its uses, and if it's treated with enough care, I think that it can play a useful role in moral debates from time to time. I hope that the argument I make in Chapter 7 is one of those times, though I'm painfully aware that I've left an important hostage to fortune by using it.
**7** Wittgenstein, 2001, § 293.
**8** Nussbaum, 2006, esp. § 3.x.
**9** Aristotle, 1992, p. 59.
**10** cf. Nussbaum, 2006, p. 181.
**11** Beckett, 2006, p. 51.
**12** Actually, he'd have a reason to prefer an attenuated future for the sake of avoiding boredom.
**13** Williams, 1973, p. 113.
**14** And more than that, actually: if infertility is a problem for Brian, it must be so not only because of an appeal to genetic relatedness, but also because it is important to him that the gametes are produced by his own body rather than the body of his identical twin – since that sibling would be capable of producing a child with exactly the same genetic relationship to him.
**15** Carel, 2011, p. 26.
**16** Ibid., p. 72.
**17** Marín-Gutzke and Sánchez-Olaso, 2010, *passim*; Ferguson and Chang, 2010, pp. 435, 444.
**18** Aristotle, 1976, p. 118.
**19** Ibid., p. 122.
**20** Vide Carel, 2009, § 5 for a tour of the literature on this.
**21** Mill, 1998, § 2.

**22** Several science fiction writers have imagined the possibility of uploading our consciousness into some electronic format; imagine, future generations might exclaim, not being able to do that, and so having to confront death, illness, non-rebootability, and all the rest of it! How did people in the early third millennium cope?
**23** Smith, 1981, pp. 869–70; emphasis added.
**24** Nussbaum, 1996, p. 21.
**25** Everett, 2009, p. 58.
**26** Harris, 2007, p. 44.
**27** Ibid., p. 36.
**28** In so doing, it might make the difference to whether or not a whole life attains a minimally acceptable goodness – though this is not decisive: there are some procedures that will not make a life reasonably acceptably good on their own, but still count as therapeutic; and at other times, a life taken as a whole might count as good enough, even though some aspect requires improvement to satisfy that demand – after all, we don't normally think that a person's life as a whole stops being acceptably good just because he has eczema: we just think that there is one part of his life that could stand to be better.
**29** Ibid., pp. 46, 56.
**30** Chadwick, 2008, p. 28.
**31** Bostrom and Ord, 2006, p. 672.

# 2

# Bad Arguments against Better Lives

The tone that I'm going to be adopting from the next chapter onwards will be one of friendly scepticism about biotechnology's role in the pursuit of the good life, laced with something more along the lines of traditional hostility to the idea that we have any obligation to develop or make use of it; but even a hostility to aspects of biotechnological evangelism comes a long way short of amounting to the idea that there is anything particularly wrong in principle with seeking to enhance or modify people. Yet there has been a number of arguments advanced that seek to defend precisely this position: that there *is* something morally problematic about many of the uses to which the biosciences might be put – particularly in respect of enhancement. The object of this chapter is to clear anti-enhancement arguments out of the way.

There is a number of forms that an anti-enhancement argument might take, but they do appear to fall into a more limited range of kinds. One of the simplest kinds of anti-enhancement argument rests on the idea that enhancement represents interference with nature, that we ought not to interfere with nature, and so ought to be at least cautious about enhancement. Nature, on this account, is granted a kind of integrity. This integrity is claimed to demand more than our respect, since respect is compatible with significant alteration; the Ashmolean museum in Oxford, for example, reopened in late 2009 after a period of substantial alteration that does not indicate any lack of respect for the integrity of the building. It is claimed to demand reverence as well.

Still: simply to point out that something is natural and therefore is somehow sacrosanct seems all-too-quickly to open the way to a hostility not only to enhancement (and remember that that hostility hasn't yet been justified), but to therapy as well. After all, vulnerability to measles, or a congenital heart problem, is perfectly natural; yet it would be odd to complain about reversing them. It would be similarly odd to complain about using spectacles to reverse age-related presbyopia, even though it is a perfectly natural process that doesn't even involve invasion by a hostile microorganism. (It's with this kind of argument that John Harris makes short work of the 'normal species functioning' account of the T/E distinction.[1]) More absurdly still, it would seem to indicate a commitment to the wrongness of dentistry, the internet, and turning on the heating in winter. Since therapy is acceptable, though, we can't hold on to the claim that interference with nature *per se* is wrong. Appeals to

nature must, therefore, be much more subtle. As we'll see, though, this line won't get us far either – especially in relation to 'human nature', given the very plausible claim that, if there is such a thing as human nature, it's that we are an inescapably unnatural species.

Habermas' appeal to human nature is different in its emphasis – his concern is not that there is anything essential about humans as a species that merits reverence, but that we have an indirect reason not to interfere in our natures because of the ethically detrimental effects that such interference might have on our social interactions. In making this argument, he provides a kind of bridge between nature arguments and dignity arguments, which I shall also examine. Finally, I'll look at Sandel's objections to enhancement, which provide a good example of arguments that revolve around the kind of attitude to the world around us, ourselves, and our place in the world that it is desirable or virtuous to cultivate.

None of these arguments will, in the end, tell us wrong with enhancement in general and in principle. Though it would be foolhardy to say that this means that there's *no* in-principle argument against it, I'm happy enough to accept that there is at least no obvious reason to think enhancing wrong in itself.[2] Still: at least these arguments *are* recognizable as arguments. This is not the case for all objections; and I'll turn to a famous non-argument against enhancement now, if only to get it out of the way.

## Repugnance as a moral tool

Bioscience, especially when it's applied to human lives, can have a tendency to attract suspicion. Terms like 'bioengineering', and the idea that we might interfere in the course of a person's development by altering them – sometimes at a cellular or genetic level – have the capacity to generate a fairly visceral reaction against them. Leon Kass is, famously, a proponent of the idea that there's a value in such reactions, and that they ought to be taken seriously in bioethical debate. He characterizes what he calls the 'wisdom of repugnance' as

> the emotional expression of deep wisdom, beyond reason's power fully to articulate it. Can anyone really give an argument fully adequate to the horror which is father-daughter incest (even with consent), or having sex with animals, or mutilating a corpse, or eating human flesh, or even just (just!) raping or murdering another human being? Would anybody's failure to give full rational justification for his or her revulsion at these practices make that revulsion ethically suspect? Not at all. On the contrary, we are suspicious of those who think that they can rationalize away our horror, say, by trying to explain the enormity of incest with arguments only about the genetic risks of inbreeding.[3]

Though Kass is talking about human cloning in this particular case, appeals to the supposed 'wisdom of repugnance' could be made in any number of cases

(Aubrey de Grey has noted that it could be applied to research into extreme longevity[4]), or in respect of biotechnology in general. The nub of the appeal is that a gut feeling can tell us something of moral import in a manner that may not be easily articulable.

The final sentence of the quoted passage is particularly revealing, nodding as it does towards the fact that, if someone is repelled by some course of action, his reactions will not necessarily be allayed by pointing out either that the risk of harm is minimal, or even by pointing out possible benefits that may be conferred by it. Pointing out that contraception was used, and that in this instance both parties had a great time, isn't likely to undermine a viscerally hostile response to hearing about incest; and much the same applies, *mutatis mutandis*, when it comes to similar reactions against using the biosciences or their products. What counts is that the very idea of a certain activity or procedure is irreducibly contaminated; the question of acceptability is not an open one.

However, and as Kass himself admits in the paragraph from which the quotation above is taken, the big problem with this kind of appeal is that it is not an argument.[5] Hence to try to present arguments against it in the manner that one would present counterarguments to concerns of the type we'll consider later in the chapter would be to mistake the essence of the rejection.

For a statement to be of *moral* interest, the speaker has to be able at the very least to be able to say that, and why, anyone who does not share his view is mistaken.[6] Clearly, this is hard if that repugnance resists or does not require articulation to begin with. Repugnance gives the person who feels it a reason to avoid some thing or practice; but it is not in itself a *moral* reason: at most, it's the stuff from which we might try to craft a moral reason at a later point – and it doesn't give a reason to anyone else.

Hence a statement along the lines that 'This is repugnant' is of little more than anthropological interest to any listener: 'Oh. It doesn't seem so to me', he might reply, without being in the least bit moved to share that repugnance. Moreover, neither will a statement of that sort indicate any obligation for the listener to change his behaviour nor do anything to halt whatever the speaker finds repugnant unless he thinks that one ought always to avoid doing things that others find repugnant; but this is as much as to admit that repugnance is trumps, and there is no non-question-begging reason at all to accept that. (To take, for example, Kass' example of anthropophagy: it might be that eating human flesh is held repugnant by everyone; but in an emergency situation, if it makes the difference between death and survival, there is no particular vice in suspending one's sense of repugnance. And this serves to illustrate that, even if repugnance is morally important, it is much more likely to be one of a range of moral considerations than it is to be the only one.)

The speaker could, perhaps, attempt to show the listener why he should be repelled by whatever it is – except that this gives the game away, by moving

from a bare statement of repugnance to something more in line with an appeal to reason and/or reasons. The most the repelled person could hope for without giving the game away is to be able to persuade the listener that his own repugnance ought to count in this case. Maybe the listener could be persuaded of that, perhaps for the sake of social harmony. But that would be a decision concerning the virtues of social harmony, rather than the allegedly repugnant practice. The long and short of it is that appeals to repugnance provide a *ground* for certain kinds of subsequent argument; but they are not sustained by it, and the argument that is sustained concerns civility far more directly than it concerns the rights and wrongs of enhancement.

Overall, I don't think I need to spend any more time on appeals to repugnance. Trying to argue against them is a fool's errand, because they aren't arguments at the core; but – for precisely that reason – neither are they likely to persuade those who aren't inclined to be repelled.

## Nature and human nature

A better way of establishing the wrongness of at least some uses of the biosciences is to construct an argument based on an appeal to nature. Ultimately, these arguments are not compelling, but they are at least arguments, following roughly the same structure: we ought not to interfere with nature (or human nature); this procedure *p* is such an interference; therefore we ought not to *p*. The minor premise, that *p* interferes with nature, is generally uncontroversial – indeed, proponents of bioscientific interventions of one sort or another frequently take it as a badge of honour that they are interfering with nature of one sort or another. Thus the focus of the debate must be on the major premise: that we ought not to interfere with nature. Bioconservatives will assent to this; others are less likely to.

The difficulty that is quickly encountered by bioconservatives is that rather quickly they end up insisting that many quite trivial things are morally problematic. I am awoken by a clock-radio, which relays voices almost instantly from locations all over the world, at exactly the same time each day; I turn on the light, and take a gulp of water from the glass by my bed. All these things could be construed as being possible only because of interference with nature – and I'm still under the quilt!

A slight variation on the same theme comes in the form of arguments along the lines that this or that is unnatural, and therefore wrong. But, as Mill pointed out, if 'nature' is understood as referring to 'the powers and properties of all things', it would be hard to imagine that anything could be unnatural:

> But in this signification, there is no need of a recommendation to act according to nature, since it is what nobody can possibly help doing, and equally whether he acts well or ill.[7]

Meanwhile, if 'nature' is supposed to indicate 'that which takes place without human intervention',

> it is evident at once that the maxim, taken in this sense, is not merely, as it is in the other sense, superfluous and unmeaning, but palpably absurd and self-contradictory. For while human action cannot help conforming to nature in the one meaning of the term, the very aim and object of action is to alter and improve Nature in the other.[8]

On one horn of the dilemma, we have to admit that *nothing* humans do is unnatural; on the other, we have to admit that the pursuit of that which nature doesn't yield is the whole point of any action in the first place – and it's not easy to see that there's anything objectionable to the mere fact of having acted. Making a fire where there previously was none is, in a sense, an unnatural act; but since there is frequently nothing objectionable at all about making a fire, it cannot be morally problematic simply by virtue of its 'unnaturalness': we need something else to say apart from an appeal to nature.

So looking to nature in the broadest sense is not likely to furnish us with all that strong an argument against the permissibility of the use of the fruits of bioscience. But if nature *qua* nature won't meet the demands of the argument, it's still possible that appealing to the importance of *aspects* of nature will be a better bet. Since this book is concerned by the contribution that the biosciences might make to the good life for humans, it makes sense therefore to look to human nature for the basis of an argument – to wonder whether there's anything about that that commands respect, and interference in which would be morally problematic. It might be, for example, that there is nothing particularly problematic about interfering with nature in the broadest terms – so the extinction of smallpox or a new treatment for diabetes would still be welcome – but that there is something essential about humanity would be held to be sacrosanct.

Ostensibly, the first problem with this kind of argument is establishing that there is such a thing as human nature, discrete from the rest of creation, to begin with. Arguably, if there is not, then we can protest all we like about its importance – it will not get us anywhere of any practical import. However, I said in Chapter 1 that I wasn't too concerned that we should be able to give precise, locked-down meanings to words; it's enough that we can recognize a human when we see one. It seems to me that any argument against the use of certain biotechnologies that relies on a particularly tight understanding of 'human nature' is likely to be weak, but unnecessarily so: there is no need to confine ourselves to any such understanding. (In the same sort of way, one doesn't have to be able to give a precise definition of 'sport' to be able to say

confidently that one doesn't like it, or that it takes up too much of the TV schedules, or whatever.)

So we can afford to allow that there is such a thing as human nature, even if we can't quite pin down for the time being what it is. Where does this take us, though? Why should we think it sacrosanct? That is not obvious. More interestingly, it turns out that it is possible to construct a plausible account of human nature that is not only capable of accommodating interference, but which has it at its heart.

In *The Artificial Ape*, Timothy Taylor advances the hypothesis that humanity is, in some ways, a profoundly enhanced, augmented creature – naturally unnatural, as it were. The nub of his claim is that it is very hard to see how a species such as *H. sapiens* could have emerged and flourished had it not been for a number of technological innovations in our ancestor species. Merely being bipedal generates all manner of problems in respect of food, for one thing. The body shape required for bipedalism means that the abdominal cavity has to shrink; this means that there is less room for intestine, and that the process of digestion is therefore less efficient. It is much less easy for a biped to get the nutrition it needs from the same amount of food than a quadruped. The solution is cooking: cooking means that we can start breaking down food before it enters the body. Cooking implements are, in effect, an artificial second stomach that our ancestors learnt to build. It was by virtue of the use of this technology that a creature like a human, which otherwise would have struggled to obtain all the calories it required, could flourish – especially given the size of its power-hungry brain.

Our brains caused us other problems, too. Bipedality requires a pelvis of a certain shape, narrower than in quadrupeds. But this means a narrower birth canal. In animals with small brains, a narrowed birth canal would be bad enough; but humans have big brains, which have to be housed in a big head. So the question presents itself: how could a creature evolve that was both big-brained and bipedal? Wouldn't these characteristics be antagonistic? The answer is that humans are born, in effect, premature: we have a massively extended infancy in which we are fairly helpless – we can't do so much as support the weight of our own head when we're born, and can't even really stand up until well into the first year. But this simply raises another problem: wouldn't such a long period of utter helplessness count against the survival both of the infant and the adults burdened with the task of raising it? Again, this makes it hard to see how a species such as *H. sapiens* could have had much chance of survival, let alone flourished: compared to other members of the animal kingdom, juvenile humans are just broken monkeys.

The solution to this puzzle, Taylor hypothesizes, can be found in the invention of tools: slings and papooses in particular. In effect, Taylor argues, our prehuman ancestors turned themselves into marsupials. Infants could then be born 'premature' and helpless, and be allowed to develop their big, brain-

filled heads *ex vivo*: the way was cleared for us to become the species we became. There is little to no direct evidence of these tools, because they would have been made from resources ready-to-hand – leaves, branches, and so on – but they do offer a simple explanation of how we became what we are.

What's important for the argument here, though, is that we can tell a story about humans in which it is central to our nature that we are the product of interference with nature. In short, humans are a technologically-formed species: we could not exist were it not for the sake of this augmentation. Thus to appeal to something like human nature as a reason not to pursue further technological alteration may not be all that promising a strategy. Admittedly, there is arguably a difference between us making alterations to ourselves with particular intentions, and pre-humans blindly doing things that would *per accidens* allow for the evolution of humans. But that would still mean that we'd have to admit that interference with human nature is generally undesirable, even if it's not always blameable; and if we accept that human nature is the product of interference, we'd have to accept that it's undesirable that humans evolved in the first place. I do not consider such a position to require a great deal to be said about it – certainly not in a book that's considering the good life. Equally, we could, I suppose, say 'This far and no further': but such lines in the sand would appear to be arbitrarily drawn – and noone has ever had and difficulty crossing a line when it's only drawn in sand. What remains much more constant than any such line is the thought that the incorporation of technology into our lives and bodies has been a characteristic of humanity for as long as there has been even the possibility of humanity.

## Habermas' future

A concern with human nature motivates Jürgen Habermas' excursion into bioethics in *The Future of Human Nature*. Though my project here is mainly concerned with things that we do to ourselves and Habermas' is mainly concerned with things that we might to do others – embryos, primarily – it is worth some consideration, if only to clear the way of as many possible anti-enhancement arguments as possible. Habermas is worried about what the possibility of enhancement technologies might to do humanity's 'ethical self-understanding', and the potential that human relationships will be distorted by them. 'Bodies stuffed with prostheses to boost performance', he suggests,

> or the intelligence of angels available on hard drives, are fantastic images. They dissolve boundaries and break connections that in our everyday actions have up to now seemed to be of an almost transcendental necessity. There is fusion of the organically grown with the technologically made, on the one hand, and separation of the productivity of the human mind from live subjectivity, on the other hand. Whether these speculations are manifestations of a feverish imagination or serious

> predictions, an expression of displaced eschatological needs or a new variety of science-fiction science, I refer to them only as examples of an instrumentalisation of human nature initiating a change in the ethical self-understanding of the species – a self-understanding no longer consistent with the normative self-understanding of persons who live in the mode of self-determination and responsible action.[9]

I'll take the speculations of the first half of this paragraph as being wholly compatible with the use of the biosciences in an endeavour to maximize flourishing. What we know now about ourselves, and what we can do to ourselves individually and as a species, does disrupt intuitions that may have been passed down for centuries. A comfortable belief in human exceptionalism has been on the rack since Darwin and Mendel, disquieting as this may have been (Nietzsche's 'madman' scene in *The Gay Science*[10] is an early articulation of the question about what we are and what we are to be in the absence of a clear, God-given framework). And knowing more about what we are has led to an awareness that what we are and what we must be are not the same. This is something that can be treated precisely as a means to increase our flourishing and the goodness of our lives, and as something thereby that we can happily embrace; and it is something also that (pace the concerns of Michael Sandel, which I address below) we can see as a liberation from the constraints of brute biology – Tuija Takala, for example, embraces the possibility of ectogenesis on the grounds that, with its development, all humans 'could finally be freed from the biologically determined roles that have for long dragged behind the social evolution'[11] – and so, on the grounds that a free life is better than an unfree one, a contribution to the good life by less direct means as well.

The sentiment of the second half of this quotation is an abiding theme of Habermas' book – the idea he develops is that a person who discovers herself to have been selected in some way may find it more difficult, if not impossible, to adopt the attitude to the world around her of a free being: the sense that she is a morally free being guided by norms and reasons may be eclipsed by the sense that she is made, rather than grown. For all humans must take an attitude to their own contingency, and all are in a sense constrained by their genome; but whereas this is not a threat to the freedom of the 'natural' human (Habermas takes seriously the idea that '[w]e experience our own freedom with reference to something which, by its very nature, is not at our disposal'[12], and that 'natural fate seems to be essential for our awareness of freedom'[13]), the same apparently cannot be said of the engineered human, who would be 'barred from developing, in the course of a reflectively appropriated and deliberately continued life history, an attitude toward her talents (and handicaps) which implies a revised self-understanding and allows for a *productive* response to the initial situation.'[14] That is, while Hamlet admits that we may at least roughhew the ends shaped for us by a divinity, Habermas worries that this would not be possible for the engineered human.

Why, though? Habermas' worry is that a child is excluded from the possibility of being an ethically equal player in the process of self-formation to the extent that she is designed. Whatever the contingencies of a natural life, each of us is equal inasmuch as that we are contingent in the same sort of way. When one person helps design another, though, this pattern is distorted:

> Up to now, only persons born, not persons made, have participated in social interaction. In the biopolitical future prophesied by liberal eugenicists, this horizontal connection would be superseded by an intergenerational stream of action and communication cutting vertically across the deliberately modified genome of future generations.[15]

So while '[a]ll persons, including those born naturally, are in one way or another dependent on their genetic program'[16],

> [o]ur concern with programming here is not whether it will restrict another person's ethical freedom and capacity of being himself, but whether, and how, it might eventually preclude a symmetrical relationship between the programmer and the product thus "designed". Eugenic programming establishes a permanent dependence between persons who know that one of them is principally barred from changing *social* places with the other.[17]

For this reason, Habermas thinks that engineering ought to be reserved to certain therapeutic cases.

Habermas' worries, if well grounded, might well inform a quest for 'the good life', too: if a person is ethically restricted in this way, then they would presumably be likely to flourish less inasmuch as that humans are creatures with certain moral capacities. But the problem for Habermas is that it's not at all clear that we should take his worries about this as being worries about enhancement technologies; rather, they're probably best served as a warning about the ethics of childrearing in general, and the moral reasons we have not to impose expectations on our offspring. There is no *prima facie* reason to think that an enhanced child would have any more difficulty engaging with her origin than would any other child; she has a decent shot at having less difficulty on this front than would be faced by a naturally formed child with overbearing parents – and being naturally born wouldn't mitigate that overbearingness, any more than being engineered militates against good parenting. His worries, in fact, tend to assume a hyperbolic filial obsequiousness, as though the discovery that one was in any way modified prenatally would dispose one to treat oneself as less morally free, and less ethically equal to others. It's not at all clear why that should be – why a person can't respond to this aspect of their pre-biography as they would to the discovery that they were conceived for any other reason, or no reason at all. Being born with the slow-twitch responses of a middle-distance runner doesn't stop anyone becoming a baker instead, or even training for the sprint.

Even if – *arguendo* – Habermas is right on all this, it's odd that his concerns don't translate to therapeutic manipulation:

> A preventively "healed" patient may later, as a person, assume a different attitude toward this type of prenatal intervention than someone who learns that his genetic makeup was programmed without his virtual consent, so to speak, according to the sole preferences of a third person. Only in the latter case does genetic intervention take on the form of an instrumentalization of human nature. In contrast to clinical intervention, the genetic material is, in this case, manipulated from the perspective of a person engaging in instrumental action in order to "collaboratively" induce, in the realm of objects, a state that is desirable according to her own goals. Genetic interventions involving the manipulation of traits constitute positive eugenics if they cross the line defined by the logic of healing, that is, the prevention of evils which one may assume to be subject to general consent.[18]

Quite how a pre-person can be treated as being able to give consent for one thing, and not for another, is never quite explained.

There is one area, to be fair, in which a Habermas-inspired suspicion of enhancement might have more power, and that is in relation to so-called 'moral enhancement'. This is because, in that context, a person's moral attitude to himself and the world around him is precisely what is manipulated according to others' desires. I'll come back to this in Chapter 7.

## The argument from dignity

One of the most common strategies deployed against the use of enhancement technologies involves the claim that they are corrosive to, or potentially corrosive to, human dignity. Article 1 of the European Union's Convention on Human Rights and Biomedicine urges that

> Parties to this Convention shall protect the dignity and identity of all human beings and guarantee everyone, without discrimination, respect for their integrity and other rights and fundamental freedoms with regard to the application of biology and medicine.

This statement comes in the context of the Preamble's claim that the signatory States are

> Convinced of the need to respect the human being both as an individual and as a member of the human species and recognising the importance of ensuring the dignity of the human being;
>
> Conscious that the misuse of biology and medicine may lead to acts endangering human dignity

and that they resolve

> to take such measures as are necessary to safeguard human dignity and the fundamental rights and freedoms of the individual with regard to the application of biology and medicine.

It is of some note that the Convention talks both about the 'dignity of human beings', and 'human dignity'. These would appear to be two different things.

The former indicates some characteristic a capacity for which is found in all token members of the species, perhaps related to decorum, and the ability to behave in a manner befitting one's place. The hero unjustly sent to the guillotine might be said to have kept his dignity to the last in this sense – and, in the process, demonstrated just what nobility of character it was that made him a hero in the first place. By contrast, the jilted lover who sells the details of an affair to the tabloids might be said to have behaved in an undignified manner – meaning that we expect something different. In the latter sense, 'dignity' indicates a quality of another kind (and, perhaps, a quality that makes individual dignity possible): something captured in the idea that, as Fukuyama puts it, 'there is something unique about the human race that entitles every member of the species to a higher moral status than the rest of the world'[19]. Slavery could provide us with an example of the sort of thing that represents an assault on dignity in this sense, inasmuch as that it denies the presence of that characteristic in the slave, treating him instead as just another thing in the world. (On this basis, we might think that even a slave who is well treated in every empirical respect still has his dignity violated simply by virtue of being owned.)

Dignity in the first sense – the dignity of individuals – would not obviously be affected by bioscientifically-informed manipulations in any way; at the most, they'd just provide another way in which a person could behave in an undignified manner, and so would be qualitatively no different from the webcam, megaphone, or quill pen. Moreover, as Alasdair Cochrane points out, even if certain modes of behaviour are more dignified than others, that still won't give us any clue about whether they're mandatory.[20] If the biosciences really do make any difference to dignity, it must be in the latter sense. Nor is it hard to see how the worries might take root. If we're in the business of altering what humans are, and human dignity is closely tied to what humans are, that dignity might seem vulnerable. The manufacture of Deltas in *Brave New World* could, on this basis, be presented as an (admittedly slightly lurid) example of how science might violate of their human dignity:

> "Consider the horse."
>
> They considered it.
>
> Mature at six; the elephant at ten. While at thirteen a man is not yet sexually mature; and is only fully grown at twenty. Hence, of course, that fruit of delayed development, the human intelligence.
>
> "But in Epsilons," said Mr Foster very justly, "we don't need human intelligence."[21]

Treatment like this invites the response that they *do* need human intelligence, because without it they would be less human – and to make a human less human would be to besmirch his status *as* human: to assault his dignity. Fukuyama articulates this worry as follows:

> What is it that we want to protect from any future advances in biotechnology? The answer is, we want to protect the full range of our complex, evolved natures against attempts at self-modification. We do not want to disrupt either the unity or the continuity of human nature, and thereby the human rights that are based on it.[22]

If we think that humanity has a particularly high moral status related to its dignity, then we might quite easily reach the conclusion that eroding that dignity is a serious wrong.

Yet the problems with this kind of claim are several. The most obvious one is that it seems to require that we have a clear idea of what, precisely, dignity is: it doesn't seem enough to say vaguely that it's simply whatever gives humans their special status; we would seem to need to know what we ought not to be eroding. Still, on the basis of the suggestion in the last chapter, I don't think that this concern needs to detain us all that much; the word 'dignity' does play a role in moral debate, and though its precise definition might be a matter of ongoing disagreement, this is not a warrant for making it an un-word quite yet. (Cochrane, incidentally, does suggest that there is nothing that the word brings to the table that we could not have otherwise, and so there would be at the very worst no loss, and probably some profit, associated with un-wording it.[23] This might be so, but it's a debate for another day.)

But there is a problem that's less easy to dismiss, which is that Fukuyama is articulating an undue pessimism about science and its products. For one thing, it is increasingly hard to sustain the idea that humans *are* radically different sorts of creatures from all the others, and the corresponding idea that humanity is unique and at one remove from the rest of creation. But, still, humans have an internal and intellectual richness that is – as far as we know – unparalleled, and so it's not completely impossible that we have a moral status that is unique; maybe human dignity can be linked to that kind of distinction. What's less clear is how that would be imperilled by the application of the biosciences. After all, the ability to manipulate ourselves even at a cellular level might very easily be thought to be coeval with whatever else makes us special; if the core of human dignity claims is that we are not just brute material, and more than the sum of the ambulant biochemistry that forms us, then it is not entirely clear why that dignity is imperilled by attempts to make ourselves even more than that.

The basic idea is that the acquisition of dignity might be the point at which a species passes through a kind of moral 'glass ceiling'; in manipulating ourselves and bodies in pursuit of the good life, it doesn't follow that we'd be stripping ourselves of any dignity, so much as outstripping ourselves.

One possible problem with this is that it might conceivably lead to a world in which a chasm opens between the biotechnological haves and have-nots; but this looks to be a problem about justice or solidarity, rather than about dignity *per se*. Sort out the justice problem, and the dignity problem will evaporate.

A further potential problem, articulated by Nicholas Agar, is that enhanced 'post-persons' may pose a threat to 'mere persons' not simply because members of the latter group would no longer be assuredly top of the moral tree, but because post persons may have a higher moral status, and so permissibly treat mere persons as a resource in the way that mere persons today treat sentient non-persons like cows or monkeys as a resource. We should not, Agar thinks, create post-persons, out of a concern for preserving our morally special status.[24] This worry I take to be overblown. Even if we think that post-persons would have a higher moral status than persons, and that higher moral status implies more entitlements, it doesn't follow that post-persons are entitled to see mere persons as expendable (although they might); to have *some* extra entitlements doesn't mean that there's nothing to which a creature isn't entitled. I could still be that, once a creature has passed what Alan Buchanan terms a threshold moral status, it has certain basic claims that are inviolable even by gods.[25] (Indeed, there's no shortage of people who think that, even though we *do* use non-persons for food and research, we have no entitlement to do so; and this is compatible with thinking that the moral status of persons is higher.) This is a point to which Michael Hauskeller and David Wasserman both correctly nod.[26]

All in all, I do not think that arguments from dignity carry all that much weight. Even if we think that humanity as a species has a status that is properly thought of in terms of dignity, the prospect of biomedical manipulation of ourselves and others is quite plausibly under threat only if we set about the task of creating Huxleyan Gammas, Deltas and Epsilons. Bioscientific intervention may provide us with new and creative ways to act in an undignified manner – but, as such, they are not qualitatively different from the technologies we already have, or consume on a night out. Even if some interventions might strike us as ill-advised, it doesn't follow that they're wrong. As Cochrane hints, it's not obvious that appeals to dignity could *ever* do that.

## A slight reprieve?

Having said all this, I think that there is a space for appeals to dignity, as long as they are handled with caution. Appeals to dignity are, as far as I can see, far more often than not unable to sustain a claim that a certain procedure is wrong. However, they might be able to take at least a little more weight if they are cast in terms of self-respect, and put to work on behalf of an appeal to some conception of the good life.

I hinted in the last chapter that a life of pure euphoria would not count as a good life. My claim was that a good life necessarily involves entertaining at least a little discomfort. A perpetual euphoric would have little reason to attend to even his basic bodily functions: he would neither be prompted to eat by a stab of hunger, nor repelled by his own filth, since he would be euphoric either way. Yet such a life would almost certainly be rejected by observers as being desirable; they might be able to accept that there is something attractive about being that happy, and never being made any less happy by anything, but it is possible to accept this idea without relinquishing the idea that, all the same, such a life would be undesirable. And the reason why it would be undesirable would seem to have to boil down to one of self-respect: whatever else I may be, I am too *good* to be unmoved by the thought of squalor. The genius who invented the euphoria pill may assure us that, if we take it, we will be unconcerned by the state of our lives. But while this is correct, a claim about the good life implies a claim about how we, as reasonable agents, would ideally want to live our lives – and this claim is irreducibly a claim about how we *at this moment* would ideally want to live our lives; being told that our concerns will vanish once we take the euphoria pill is therefore to miss the point that it is that euphoria pill that concerns us in the first place. We might well not regret it; but we wouldn't regret that same genius killing us, either, and that's doesn't mean that it's irrational for us to want not to be killed.

What is acceptable as a part of the good life is partly shaped by self-respect; and continued self-respect is also a constituent part of ongoing flourishing. And it makes sense to treat self-respect as being more or less synonymous with dignity. Dignity, that is, will inform what we accept as a good life – for example, we might reject certain things as being *infra dig*, or perhaps contrary to our *amour de soi* – and (*qua* a certain kind of fortitude) will also appear as a part of that good life.

In Chapters 6 and 7, I'll invoke something like a dignity argument to dampen some of the enthusiasm that we might have for certain kinds of cognitive or moral enhancement. In both cases, it is possible that these 'enhancements' could be used not so much to help agents realize their visions of the good life, as to corral them into lives of a certain sort. Many enhancements are held to be desirable because they enable a person to function better in the world as they find it and to be able to meet more effectively the demands that it makes. But this sort of claim tacitly assumes that, when there is a conflict between what the world demands of a person, and what she actually can do, it is the person that ought to change. Sometimes it might be; but it is not at all clear that this will always be the case. For it always to be people that change to meet the demands of, say, a given economic system is for them to be treated as functionaries of that system. This is not obviously compatible with their living a flourishing life; there is a meaningful sense in which this is because it is an undignified life – a life that self-respect resists. Though I shall be criticizing Michael Sandel later in

the chapter, this is an idea he captures well in *The Case against Perfection*: 'It is tempting', he writes, 'to think that bioengineering our children and ourselves for success in a competitive society is an exercise if freedom.'

> But changing our nature to fit the world, rather than the other way around, is actually the deepest form of disempowerment. It distracts us from reflecting critically on the world, and deadens the impulse to social and political improvement.[27]

Of course, some might defend enhancement on the basis that there is little difference between improving the person and improving the world. But, while the two end up meshing better either way, and this is an improvement however you look at it, the defence still misfires. And this is simply because a life that is expected always to adapt to the world is one that is expected to run to stay still – an expectation that is exhausting enough in its own right, but more demanding yet when we remember that the person expected to change would have an investment in a particular kind of life that may have to be abandoned. Offering people the chance to enhance themselves as a means to catch up with the world around them is, when it comes to it, to offer them Hobson's choice.

## The mythologization of the given

Michael Sandel describes his programme in *The Case against Perfection* as 'opposing genetic enhancement'[28]. (Though he does primarily talk about *genetic* enhancement, the arguments translate to enhancement generally well enough for my purposes here.) The tenor of the essay indicates that this opposition is not dogmatic, so much as based on a claim that enhancement is incompatible with, and damaging to, something else that Sandel thinks important – something he calls an 'ethic of giftedness'. His position appears to be that there are good things to be had from enhancement, but that they're countered by the good that enhancement precludes.

The idea of the gift, giftedness, and the given is one that recurs many times throughout the essay. For Sandel,

> [t]o acknowledge the giftedness of life is to recognize that our talents and powers are not wholly our own doing, nor even fully ours, despite the efforts we expend to develop and to exercise them. It is also to recognize that not everything in the world is open to any use that we may desire or devise. An appreciation of the giftedness of life constrains the Promethean project and conduces to a certain humility.[29]

A little later, he claims that

> eugenic parenting is objectionable because it expresses and entrenches a certain stance toward the world – a stance of mastery and dominion that fails to appreciate

> the gifted character of human powers and achievements, and misses the part of freedom that consists in a persisting negotiation with the given.[30]

There is a need for a certain amount of untangling here, since the sense of terms like 'giftedness', 'gifted' and 'the given' is not wholly clear. And the more one thinks about the terms, the more puzzling they become. On a superficial level, Sandel's claim seems to be that there is a virtue in taking the world as one finds it, and accommodating oneself to it. That there is a *degree* of virtue in such an attitude is not wholly implausible – it might amount to the idea, for example, that it is wise to know what battles are worth fighting – and I've already conceded in the last chapter that there might be times when a person could justifiably be persuaded away from a particular vision of the good life. Yet taking the world as one finds it cannot be taken seriously as an attitude worth endorsing. Someone who refuses to use the telephone or the printed press, and explains this by saying that the virtue of taking the world as one finds it militates against communicating with anyone who is not in the same room, would be absurd. Neither does someone who refuses to take painkillers to fend off a migraine because he accepts the world as a given have virtue on his side. He would be a fool, and foolishness is a vice. As Kass notes:

> The mere "giftedness" of things cannot tell us which gifts are to be accepted as is, which are to be improved through use or training, which are to be housebroken through self-command or medication, and which opposed like the plague.[31]

So the most that Sandel can mean is that there is something to be said for taking the world as one finds it *to some extent*. And this would be captured in the idea that there is something to be said for 'negotiation' with the given. Yet this is puzzling, too: 'the given' – presumably, what a phenomenologically-minded thinker would call 'facticity' – is a poor negotiating partner, because it won't negotiate back. Sandel must, then, be speaking metaphorically, and talking again about a decision concerning the extent to which we are prepared to take the world as we find it.

Sandel's attempt to fill the gaps appears to ride on an appeal to freedom; he does, as we have seen, talk about 'the part of freedom that consists in a persisting negotiation with the given'. What, though, are we supposed to make of this? In what sense is freedom constituted by 'a persisting negotiation with the given'? In normal language, my entering into a negotiation with you would only be required when our ambitions clash, and when we have to decide how much of those ambitions we are respectively going to have to relinquish. To this extent, negotiation is a matter of setting boundaries. Negotiation might mean my being prepared to relinquish some dimension of my freedom for the sake of preserving yours, and *vice versa*, when the two come into conflict. (e.g., negotiation might mean my having to relinquish some of my freedom to play loud music, and – less plausibly – you relinquishing some of your freedom

never to overhear it at all.) Yet while one must be free in order to be able to enter into a negotiation, quite how freedom *consists in* negotiation, remains a mystery – even if we think that it's possible to negotiate with the given. What kind of freedom is this?

(Since I've already made a reference to facticity, we might be tempted to draw a parallel with Heidegger's notion of freedom in the essay 'On the Essence of Truth' as 'letting beings be'[32] – but Heidegger is even less clear than Sandel; and, anyway: being able to draw a parallel between two things isn't enough to clarify either.[33] Assessing the truth of either's claim is no small task.)

The appeal to freedom is unilluminating. Even if we think that freedom would be restricted in a world where biotechnological interventions such as enhancement were expected, this isn't what Sandel says. Is there any better way to unpack the 'ethic of giftedness' than by appealing to freedom?

There might be, but a better question concerns whether there is actually much need to – and there might very well not be. It is not completely *outré* to think that there is a virtue in finding satisfaction with the world as one finds it, even if this doesn't quite amount to Nietzschean *amor fati*. Although I suggested a little while ago that it can be pushed to absurd lengths, we call an ability to endure difficulty fortitude, and frequently treat that characteristic as admirable and worth cultivating; and, this being the case, the idea that there might be something admirable about being satisfied with what one has ought to follow as a matter of course. This is certainly the case in respect of those situations in which nothing can be done, since a refusal to accept certain aspects of the world can prevent us making the best of the life that we have (or could ever plausibly expect to have), and might even lead to making the worst of it. (To take an example to which I'll return in Chapter 4, it might be possible to spend so much time worrying about the badness of dying that it seeps into everything a person does. Trying to avoid death might turn out to be inimical to leading a good life if it means having something *ex hypothesi* undesirable constantly playing on the back of your mind. To the same extent, accepting mortality – within the bounds of reasonably prudent behaviour – may free a person to get on with the serious business of making the most of whatever time she has left.)

So we can accept that there's a wisdom, and a virtue, in knowing which battles aren't worth fighting. Still, it doesn't follow from that that there's anything suspicious about taking what means *are* available, or even extending the scope of available means, to make it otherwise. And it's not stretching the concept of freedom too far to suggest that taking those means is a liberating project. Yet this kind of liberation seems to be quite far removed from Sandel's idea of the freedom he thinks preserved in a 'negotiation with the given'. An 'ethic of giftedness', understood in this sense, does not need much in the way of unpacking: it just means that it is both wise and admirable to keep a sense of perspective about one's life and possibilities.

Importantly, though, what it does *not* amount to is a dogma of giftedness – one can accept the world as a given without having to accept its immutability; it makes sense to imagine someone conducting a kind of audit of the world and his life, and saying to himself, 'Well, this is how things are; what now?' There is no reason to think that there is an *a priori* upper limit to the amount of modification and enhancement that one may undergo; it means at most that there is something rather foolish about pursuing them without much in the way of thought and consideration. This is a rather trivial statement, though; and it is almost completely non-directive.

It is also not at all clear why an attitude of mastery and domination over the world is all that problematic. It's possible that what Sandel has in mind here is an echo of Heideggerian worries about 'enframing' the world and humanity – see, for example, his essay 'The Question Concerning Technology'[34] – but Heidegger's position requires a great deal of preliminary work that Sandel does not appeal to countenance, and is, in any case, not obviously correct and not obviously a moral claim anyway.[35] If we do want to moralize an *attitude* of domination, we have to be able to say either that (a) the thing the domination of which is being willed is thereby wronged or (b) that there is a danger that a domineering attitude in this case, though not inappropriate, will lead to our taking a similar attitude in another case where it *would* be inappropriate: this is, in effect, a slippery slope argument.

There is no reason to believe that 'nature' or the given is wronged by our attempts to fit it to our wills – or how it even could be; so all that remains is the slippery slope argument. Yet slippery slope arguments are notoriously slippery in their own right, and Govert den Hartogh has provided an ample demonstration of why this is so.[36] Their power relies in no small part on the supposition that, in allowing the move from permissible action *A* to permissible action *B*, we will have made it less easy to hold back from impermissible action *N* – which means, in effect, that we will forget how to make moral judgements. Since there is no reason to suppose that we will lose any such power, the gradient and the slipperiness of the slope remain unclear.

The 'ethic of giftedness' is, in the end, a very puzzling thing; it almost certainly won't tell us that we ought to eschew enhancement for ourselves or others. Stephen Wilkinson suggests a response that is pithy and directed. Taking aim at the idea that it's wrong for parents to try to select characteristics in their future children because those children are and should be treated as a gift, he points out that

> I might receive a hideous gift, or one with exceedingly high running costs. It seems certain that there will be *some* cases where it is permissible to not to react with feelings of gratitude. Hence the child-gift analogy fails to support the claim that parents are required to feel appreciative of their children no matter what they are like.[37]

The point translates easily enough. Being gifted is one thing: being required to put up with that gift is quite another. Indeed, while we might have qualms about appropriating a gift for a purpose different from that of the donor out of a sense of deference for her feelings, we don't even have to worry about this when the 'donor' is blind, dumb, indifferent nature.

## Is enhancement permissible?

As I admitted at the start of the chapter, it's impossible to say that there could not in principle be an argument showing that enhancement is impermissible simply on the basis that one has rebutted or rejected a number of extant arguments. However, I'm equally happy to insist that, while absence of evidence is not evidence of absence, the onus is on the opponent of enhancement to show that there is something inescapably objectionable about the whole enterprise; *sans* such a demonstration, it is perfectly legitimate to work on the assumption that there is no such argument. And on that basis, I shall now leave bioconservative arguments for enhancement's impermissibility to one side, and turn my attention to its antithesis: the idea that enhancement is obligatory. As we'll see, it doesn't fare much better.

## Notes

**1** Harris, 2007, *passim*.
**2** Appeals to dignity will get a look in during Chapter 7, but the context will be different, and even there, I don't pretend that they'll apply to *all* enhancement.
**3** Kass, 1997, p. 687.
**4** see de Grey 2004 and 2005, *passim*.
**5** Kass, 1997, p. 687.
**6** For more on this theme, see Brassington, 2013, *passim*.
**7** Mill, 1874, p. 15.
**8** Ibid., p. 19.
**9** Habermas, 2008, p. 42.
**10** Nietzsche, 1974, § 125.
**11** Takala, 2009, p. 194.
**12** Habermas, 2008, p. 58.
**13** Ibid., p. 60; slightly modified.
**14** Ibid., p. 62.
**15** Ibid., p. 65; cf. p. 42.
**16** Ibid., p. 64.
**17** Ibid., p. 65.
**18** Ibid., p. 52.

**19** Fukuyama, 2002, p. 160.
**20** Cochrane, 2010, p. 235.
**21** Huxley, 1977, p. 30.
**22** Fukuyama, 2002, p. 172.
**23** Cochrane, 2010, p. 241.
**24** Agar, 2013, *passim*.
**25** Buchanan, 2009, p. 367.
**26** Hauskeller, 2013, p. 76; Wasserman, 2013, p. 79.
**27** Sandel, 2007, p. 97.
**28** Ibid., p. 101.
**29** Ibid., p. 27.
**30** Ibid., p. 83.
**31** Kass, 2003, p. 19.
**32** repr. in Heidegger, 1999.
**33** Thiele (1994, p. 284), for example, tries to elaborate Heidegger's point by describing freedom as allowing humans 'to glance beyond . . . being and its mastery in thought, word, or deed'; but why this counts as *freedom* is no more clear.
**34** repr. in Ibid.
**35** See Brassington, 2007, *passim*, for more on the ethical aspects of Heidegger and technology as 'enframing'; see also Hodge, 1995, *passim*'
**36** den Hartogh, 2001, pp. 283–4.
**37** Wilkinson, 2010, p. 40, slightly modified.

# 3

# Must We Make Better People?

Progress in the biosciences promises all kinds of opportunities to do things to ourselves in the hope of obtaining a better life; and, as I argued in the last chapter, arguments for the impermissibility of such interventions are not compelling. At the other end of the spectrum, there is a range of arguments that suggest that it is not only permissible to make use of these technologies, but impermissible *not* to. In this chapter, I'll examine such arguments – mainly through the prism of John Harris' claim about there being a duty to enhance. Claims about duties to enhance will often refer to things that one might do to (or for) members of the next generation; not the least of the reasons for this is that a claim that we have a duty to enhance ourselves – even if you accept that there is such a thing as a duty to self, and could supply an argument to show that it would apply in this case[1] – looks to be highly paternalistic if it's supposed to be for our own sakes, and slightly totalitarian if it's supposed to be for others'.[2]

I shall assume that, in respect of the putative duty to enhance others, the obligation is based on the good that it can purportedly do. This being the case, part of the work that the proponent of obligatoriness has to do involves establishing that beneficence is a duty. I'll consider this matter briefly in this chapter, and again (at some length) in Chapter 8, with the intention of arguing that it isn't. However, even without considering that, I think that it can be shown that there is little to no reason to believe that enhancement is a moral duty. On the other hand, my arguments about a duty to enhance have no bearing on whether or not there's a duty to provide therapy. To iterate the argument articulated in Chapter 1, though there is a sense in which therapy is an enhancement, and that might lead us to wonder whether, if enhancement is not a duty, therapy might therefore also not be a duty, 'enhancement' in the everyday sense means something different; and it's that everyday sense that I have in mind here. I'm going to accept without argument that providing therapy, in the sense of providing relief from suffering or aiming to ensure that a life is at least of a reasonably acceptable quality, is obligatory – or, at the very least, quite likely to be a duty.

## John Harris' argument for a duty to enhance

There is a brand of optimism insisting that biomedical technologies – particularly those that involve genetic manipulation – promise to allow humanity to achieve in a generation or two what might take natural selection millennia in respect of the elimination of certain diseases or tendencies to disease, and increase in desirable traits such as intelligence. It might even be suggested that those same technologies could be used to instil desirable characteristics that we might reasonably expect *never* to appear 'naturally'. John Harris belongs to this optimistic group; he claims in his book *Enhancing Evolution* that we have a duty to make use of enhancement techniques to improve the lives of future generations, and uses the word 'enhancement' to classify any such procedure. As we'll see, and notwithstanding the arguments I've already presented, the use of the word 'enhancement' in this way stores up problems; but it is nevertheless true that some enhancements are possible now, and some of those that are not currently possible are not so squarely in the realm of speculation as to mean that duties relating to them are condemned always to be merely formalistic. Hence I'll take it as read for the sake of this chapter that the substantial modification of humans for their own good is either possible or theoretically possible. Whether we have a duty to research in order to make practically possible what is only theoretically possible at the moment is a question I'll address in Chapters 8 and 9.

Closely related to Harris' claims about enhancement are claims such as those advanced by people like Julian Savulescu concerning 'procreative beneficence'.[3] The principle of procreative beneficence (PPB) rests on the idea that there are traits that it is better for future children not to have. For example, asthma may be easily controllable, but it is probably better not to have it at all than to have to control it. In a situation in which a fertility doctor can implant no more than one of a pair of embryos into a woman undergoing IVF, and one of those embryos carries the 'asthma gene' (there is no such gene as far as I know, but the term is a useful shorthand), it would be better – to the point of being obligatory – to implant the one without. After all, if one of those embryos is not going to be brought to birth anyway, it is better on the whole that it should be the one expected to have the worse life, however small the difference.

Still, whatever the resemblance, Savulescu and Harris are not necessarily arguing for quite the same thing. There appears to be a couple of important differences. Most obviously, the defence of the PPB concentrates mainly on choosing which embryo to implant; Harris' claims encompass this, but also nod towards the idea of inserting characteristics that would not otherwise have been there. A second difference is that procreative beneficence claims are overwhelmingly 'negative': they concentrate on avoiding traits that a reasonable person would want (all else considered) to be without, and then getting rid of it. Even if asthma does not seriously erode flourishing, it would still be reasonable

to think that it is better not to be asthmatic. (This is presumably why asthmatics take medication – to get as close as possible to not being asthmatic.) The kind of thing for which Harris argues, by contrast, is more 'positive', inasmuch as that it is concerned to *promote* particular characteristics; it is a slightly more radical claim for this reason. And we might think that this makes a difference: while positive enhancement is more of a piece with what I called 'bare enhancement', removing characteristics that a reasonable person would not miss has a more therapeutic tinge.

But there are limits to the plausibility of a therapeutic interpretation of procreative beneficence. One reason for this – albeit the least powerful – relies on noting that, while it may be better to be rid of something like asthma for their annoyance factor, this kind of condition is frequently compatible with a more-than-minimally reasonable level of flourishing; hence the beneficent elimination of the asthma gene in a given person would not necessarily make anything more than the most marginal difference to his quality of life, and it's unlikely to be determinative of whether or not he's reasonably flourishing. It would not satisfy the criterion for therapy that I suggested in Chapter 1. Yet this attempt to limit the interpretation is weak precisely because there is no reason to think that a therapeutic intervention properly so-called would have to make a difference to the quality of a whole life; it might not. Treating the infected wound of the refugee from Chapter 1 would self-evidently count as therapeutic even though the quality of her life might still be, on the whole, unacceptably low.

A more powerful reason to think that there's a limit is just that Savulescu's thought experiment is one in which we're to choose between two embryos, one of which will certainly come to birth, and the other of which certainly won't. There is not obviously anything therapeutic about this; at most, we're minimizing the chance that there'll be a need for therapy later on. PPB is about creating the child with the best chance of the best life[4] – it's not about improving lives. As such, the argument is perhaps best thought of as being neither therapeutic nor enhancing, but prophylactic.

Whether or not there's an obligation to make use of this prophylaxis is another matter: Rebecca Bennett has argued that there is not.[5] There is no need to go into the argument in any great depth here; it suffices to say that the principle of PPB aims to generate the best possible future world based on the avoidance of impersonal harm – and it's an examination of the idea of impersonal harm that occupies Bennett. The overwhelming majority of actual lives are not passed in a 'harmed state', she reasons – that is: most lives, even those we would prefer to avoid, are worth living. But if no actual person is harmed by being brought to birth that makes it hard to see how there is a compelling moral reason to avoid bringing that person to birth. If some lives are actually not worth living, there might conceivably be an obligation not to bring them to birth – but this is a long way from establishing that there might

be an obligation to choose against the asthma gene. On this basis, we can infer that the PBB is not just about avoiding harm, but about avoiding good-enough, or even good, lives, just so long as they could counterfactually be better. This would bring us towards something more Harrissian, but would rely also on there being an obligation to be beneficent. This is something I'll dispute here and in Chapter 8. (Even those who, like Peter Herissone-Kelly,[6] have qualms about Bennett's arguments are not *de facto* committed to accepting the PBB, either.[7])

For the most part in this chapter, I think it is fair to treat Harris' argument as representative of the kinds of argument for a duty to enhance; the problems that it faces are more or less the problems that face any assertion that there is a duty to enhance; and there *are* problems. The case that I'll be making is that the hand that people like Harris have is reasonably strong, but is overplayed all the same; and it's a tendency to overplay that bedevils claims about an duty to enhance more generally. While, for the most part, there is no duty *not* to enhance, neither is there a duty to do so. The question that remains is whether enhancement is a particularly good idea short of being obligatory, and I'll confront that in the next few chapters.

## Harris' argument

The 'ethical case for making better people' is spelt out in the opening chapters of *Enhancing Evolution* (and particularly in Chapter 2, the title of which – 'Enhancement is a Moral Duty' – sets out the stall fairly clearly). A crucial part of establishing his claim is the definition of enhancement as 'anything that makes . . . a difference for the better'[8]. Were a procedure not to make a difference for the better, it would be a straightforward semantic mistake to call it an enhancement. Obviously, I have a slightly different understanding of what makes for an enhancement; but for the moment, I will let this pass.

Good things may be good in one of two respects, Harris thinks. 'Positionally' good things are good because they confer a particular comparative advantage to their holder. For example, if being tall is a good thing, it is likely to be because of the advantages that the tall have over the short. A world in which the average adult height was around 150 or around 200 centimetres would be neither better nor worse than the one in which we live now. However, a person around 175 centimetres tall might well have a significantly different balance of advantages and disadvantages in each world in terms of shelves being reachable, doorframes being headbuttable and so on. Being a certain height is good or bad depending on the height of those around you.

Other things, by contrast, may be 'absolutely' good. There is, for Harris, a set of characteristics that do not necessarily give us advantages over others (although they may), but that simply make lives better.[9] Education is an example

of the sort of thing that we might plausibly hold to be absolutely good: being better educated than one's peers is likely to confer positional benefit, but being educated is something desirable in its own right as well. A world in which only a few is educated would be a world in which that few would have a positional advantage over the uneducated; that advantage would be eroded were the hitherto-uneducated to receive an education. But benefit of being educated *at all* would be undiminished.

Harris relates the defence of enhancement to the absolute benefits it promises:

> I favour and defend enhancements as absolute rather than as positional goods. I defend them because they are good for people, not because they confer advantages on some but not on others.[10]

It is not a given that anyone would really have to limit their defence of enhancement to cover positional goods, even by Harrissian lights. After all, the claim is also advanced that 'it is not wrong to advantage some in ways that also confer a positional advantage when all cannot be bettered in those ways.'[11] Having said that, the idea that it is *mandatory* to confer positional benefit would be harder to defend (or construct to begin with) than would either the idea that it is *permissible* to do so, or the idea that it is mandatory to confer absolute benefit. (David Jensen's Kantian argument against enhancement for positional advantage is, though problematic in some ways, rather neat on this particular point.[12]) Nor would someone who is concerned about making use of enhancement technologies for the sake of the good life really have cause to worry about positional advantage (unless, implausibly, she could demonstrate that her life would be made substantially worse merely by virtue of there being someone in comparison to whom she was positionally disadvantaged). Yet whether enhancements are better defended as positional or absolute goods, or as both, is a side-issue: more central is the point that enhancements, somatic or genetic, are worth having *qua* enhancements; their value lies primarily in the fact that the bearers of those enhancements would flourish more than they otherwise might.

Having established that, Harris moves on to claim that enhancement is much less occult a thing than many would assume: spectacles, telescopes and microscopes make some kind of difference for the better, and therefore represent fairly innocuous ocular enhancement technologies (and their being able to make a difference for the better is sufficient, in this account, to show that they are enhancements); inoculation against disease is another. Alteration of the human body or genome with the intention of conferring some desired characteristic is, the argument goes, not different in kind. Someone who accepts things like contact lenses and vaccinations ought, on pain of inconsistency, to accept the principle of enhancement more generally.

(A possible interjection here is that the advocate of enhancement has conflated enhancement and equipment: the two categories do not overlap, and

while something like genetic modification belongs in the former, things like spectacles belong in the latter. There might be, on this account, a difference between improving *people*, and improving *function*. Yet it is not obvious that this interjection will take us all that far. For one thing, if we accept that an enhancement is anything that makes a difference for the better, that is all we need to be able to say that to be equipped is to be enhanced. For another, the improvements offered by something like bioengineering are likely to be assessed in terms of the difference they make to function – just as are the improvements offered by becoming better equipped. So the enhancement/equipment distinction is probably merely semantic; and if there is more to it than semantics, it still remains to be seen that it is important.)

There is no need to go further along this line. This short rehearsal is meant to demonstrate that Harris has the components of an argument for enhancement. But demonstrating the components of an argument is not the same as *making* an argument for 'making better people', and falls long way short of establishing that such is a duty.

A basic version of the missing parts of the argument is offered in Chapter 3 of *Enhancing Evolution*, where Harris claims that 'improving life, health, life-expectancy and so on is . . . a *mandatory* dimension of any moral programme'.[13] The mandate derives from the consequentialist idea that moral value is a function of welfare, so that the power of a duty derives from the benefit that it gives, and the utilitarian idea that builds on this, that we have an obligation to maximise welfare. If we are satisfied that something like genetic enhancement is capable of improving life, health, life-expectancy and so on, the natural next step is to say that genetic enhancement is thereby a duty. I shall claim that Harris mis-steps here in a little while – but there is more to say before I get that far. For now, it is enough to note that Harris' claim at this point represents at least the beginnings of a positive argument that is stated and elaborated in several places throughout *Enhancing Evolution*.

For example, in relation to the possibility that we might engineer increased longevity into ourselves or our children, Harris suggests that '[d]ying of old age is acceptable only because it is regarded as inevitable. Absent the inevitability, the acceptability is problematic.'[14] The logic of the passage would seem to be capable of bearing something stronger: if the sole condition for the acceptability of death from old age is its perceived inevitability, and were it (justifiably) to become regarded as avoidable, it would *not* be acceptable. And had we the means to get rid of something the presence of which is not acceptable, it would seem to follow that we would have a good moral reason to get rid of it – a duty, in fact, since to call something 'unacceptable' in everyday language is to say that it *ought* to be altered or got rid of. Harris is saying that allowing people to die from old age in a world in which they need not is not acceptable; the possibility of enhancement of the right sort means that we live, or may soon live, in just such a world. Others – notably, Aubrey de Grey – have mounted arguments that amount to largely the same thing.

At the same time, appealing to the normality of senescence pulls no moral weight:

> If we subtract the acceptability from the normality we are left with nothing of moral, nor of argumentative force. Traits in short are not acceptable (in the normative rather than of course the simply descriptive sense of 'acceptable') because they are normal, they are acceptable because they are worth having. If they are not worth having, or if they are worth not having, their normality seems bereft of interest or force. This, it seems to me, obviously applies to saving life, that is to postponing death to increase longevity, resistance to the diseases of old age and to those that strike at any time like HIV, cancer, and heart disease.[15]

If mortality was a characteristic worth having, then dying might be counted as unproblematic. But it is not something worth having, and dying is something that (*pace* what I'll say in the next chapter) would be better not to happen.[16] And if it is better that something not happen, then – granted the (*prima facie* minimal) premise that, where possible, we ought not to allow that which it would be better not to happen – this generates the claim that we ought not to allow avoidable death.

To reinforce the argument, note the close association that Harris draws in the passage quoted between saving a life and postponing death or increasing longevity. For him, to postpone a death just is to save a life: the phrases are interchangeable and, therefore, their moral weight must be the same. For Harris, there is no difference between preventing the harm of death from age, Aids or cancer and conferring the benefit of extra life or health. Moreover, since preventing these deaths is a matter of saving lives, it is something we ought to treat as a duty, in just the same way as we are duty-bound to save, or to attempt to save, a drowning child.

If this goes through, Harris thinks, it elides the difference between preventing harm and conferring benefit *tout court*, since it is the conferral of benefit that does the moral work. *Any* conferral of benefit will, by the same token, have corresponding moral weight. When it comes to leaping from claims about extra longevity or disease resistance to claims about, say, sharpened intellectual skills, '[t]here seems to be no difference in principle here and thus no difference in the relevant ethical considerations'.[17]

In a nutshell, Harris' position boils down to the claim that, where we can improve lives by whatever standard we choose to assess those lives, we ought; enhancement can improve lives; therefore we ought to enhance. Granted that enhancement may require further scientific research, Harris thinks that there is a secondary obligation to support and/or participate in that[18] – and I'll turn my attention to that claim later.

Using Harris' argument as a template, I want to spend the next few pages arguing that there is no duty to enhance. This said, I am willing to countenance the possibility that many enhancements are not only permissible, but admirable as well. But there is no obligation to be admirable; and I think that the arguments for there being an obligation to enhance are overplayed

somewhat, and a number of questions remains unanswered. The areas on which I'll mainly concentrate here concern the best understanding of the two main words in the claim that enhancement is a duty: what is meant by 'duty', and by 'enhancement'?

## Why would enhancement be a duty?

In order to get to grips with the question of whether or not enhancement is a duty, it may well be useful to spend a little time looking at what constitutes a duty to begin with.

Very often, we don't have to worry too much about the precise meaning of the word; but in the face of a radical claim that there is a duty to enhance, a little circumspection might be warranted. Harris does not offer a definition of the word 'duty' within the confines of *Enhancing Evolution* – but he has said that 'we have a moral duty to do what we have good moral reasons for doing'[19]. This can't be all there is to it, though, for a couple of reasons. For one thing, an obligation is a reason to act; but it doesn't follow from that that all reasons to act generate duties. (The victim of a crime has a reason to punch the perpetrator; but there are reasons not to, too, and there is certainly no duty to punch him.) For another, 'duty' describes not only that which we have a good reason to do, but something more: it describes actions that others could justifiably blame us for not having carried out. That is to say: a duty is an action the omission of which is blameable. The difference will turn out to be important.

Within this definition, there is scope to draw a distinction between *prima facie* and absolute duties. After Ross,[20] a duty may be classed as *prima facie* if (a) the action it describes is obligatory in the normal run of things, but (b) the proper consideration of other duties or rights may nullify that obligation. By contrast, an absolute duty is that action the omission of which would be blameable in all cases. Whether or not there is such a thing as an absolute duty, or whether blameability is always a function of context, would be a further question – but it would be a question that belongs in the field of meta-ethics, and is not one that I can address here. For the time being, it's enough to say that, *if* we're going to be making appeals to duty, a distinction between absolute duties and rebuttable *prima facie* duties is a plausible one to draw.

If enhancement is a moral duty, then it would have to follow that *not* to enhance is blameable in at least some cases. So we may justly ask whether, and under what circumstances, it is true that the failure to enhance is blameable. By raising this question, we may hope to get more of a handle on the nature of the duty.

One possibility, which can be shut down fairly quickly, is to relate duties to enhance to a duty of non-maleficence. That there is a duty of non-maleficence

is an idea familiar within bioethics from Beauchamp and Childress' attempt to capture the essence of biomedical ethics in four principles[21]; I shall take it as read that having a duty not deliberately to cause harm is trivially acceptable in its own right. Some commentators have claimed that to fail to promote the welfare of others is to harm them: for example, Harris has claimed that work on cancer therapy is justified by means of an appeal to ' "beneficence or nonmaleficence", or welfare'[22] – note that beneficence and non-maleficence are treated as describing the same basic concern – and has written that 'the withholding of a benefit that could be conferred harms the potential recipient'[23]. If enhancement confers benefit, not conferring benefit is a harm; and since harming violates the duty of non-maleficence, there would be a duty to enhance.

But while non-maleficence and beneficence are related, inasmuch as both have at least something to do with welfare claims (and so, indirectly, to do with flourishing), non-maleficence is not the same as beneficence. The former is negative in its orientation: it is concerned with avoiding the causation of added suffering. The latter is positive, being concerned with the promotion of welfare. Suffering and welfare might belong on the same spectrum, but this does not mean that non-maleficence and beneficence are versions of the same underlying disposition. I can easily enough refrain from making your life worse than it is now without making it any better: if withholding a benefit harms the potential recipient, it would seem to follow that I have harmed the people I meet by not offering them the contents of my wallet; and if harming is wrong, it would follow that (*ceteris paribus*) not offering them the contents of my wallet is wrong also. That, I take it, is absurd. Nor is my refusal to steal from you an attenuated form of charity. This ought to be enough to demonstrate that duties of beneficence and duties of non-maleficence are, and should be kept, separate.

All the same, we might be able to ground a duty to enhance in a duty of beneficence that could be established argumentatively on its own terms. If a duty to enhance obtains because enhancement improves people's lives, this implies that the real moral weight is carried by the prospect of that improvement; this implies a duty of beneficence. In this respect, I have already noted that Harris is committed to the idea that the moral power of a duty derives from the benefit that it gives: at the start of *Enhancing Evolution*, he talks about a 'responsibility shared by all moral agents to make the world a better place'.[24] I take it that that which one has a responsibility to do is equivalent to that which one has a duty to do. So Harris wants to argue that, since enhancement (by definition) makes individuals or individuals' lives better than they would otherwise be, and since a world in which a person's life is improved is better than one in which that life is not improved, enhancement is a duty. This, in turn, implies (on my understanding) that it would be blameable not to improve those lives – which is as much as to say that it would be blameable not to make

the world a better place. And, though welfare is not the same as flourishing, it is straightforward enough to say that an increase in welfare will generally correspond to an increase in flourishing. Hence I shall take the idea of a duty of beneficence to be wide enough to cover a duty to promote the flourishing of others.

The question then would be this: would it be blameable not to promote others' flourishing?

Sometimes, it might be. By improving a person's level of flourishing to a point at which it becomes reasonably acceptable, we would be promoting that flourishing. I'm happy enough to allow that making such improvements where possible is at least a *prima facie* obligation. However, by the standards of the argument set out in Chapter 1, this would count as a therapeutic intervention, rather than bare enhancement. It would not, therefore, be enough to establish that there is a more general duty to enhance. To do that, we'd have to be able to say that it's blameable not to provide for the further improvement of lives that are already at least reasonably acceptable in terms of flourishing. And while providing such improvement may well be admirable, I simply don't see that a failure to provide them would be blameable. Hence there is not a duty.

Here, there might be a worry that I am simply gainsaying or talking past the Harrissian claim, and that using a particular understanding of duty to undermine a claim that enhancement is a duty is all very well, but won't serve to persuade anyone signed up to the Harrissian account. It remains to be shown either that there is no duty of beneficence, or that it is not wide enough in its scope to mandate enhancement.

I think that this can be shown, on the grounds that it is not at all obvious how establishing a beneficence-based and generalized duty to enhance might be done. I'll go into the argument in more depth in Chapter 8, but it seems that there's a range of ways in which we might set about the task – though none of them is up to it.

The least plausible strategy would be to appeal to a version of virtue ethics, and promote a disposition to beneficence – to do the kind of thing that improves the lives of others irrespective of how well-off they are already – as the kind of quality we would admire. Still, there are hurdles that need to be jumped before we can say anything about a duty to enhance from this starting point. Notably, however welcome a moral trait such as a disposition to beneficence would be, strong *ought* statements fit only uncomfortably within a virtue-ethical framework. Hence claims about duty tend to be eclipsed by claims about admirability – but while few, if any, would deny that beneficence is admirable, that won't establish obligation without adverting to some non-virtue-ethical framework; and if we're prepared to step outside of that framework, it stops being obvious why we stepped into it to begin with, or whether it was even a good move.

Various deontological theories tell us that there's a duty to assist the needy – Kant, for example, thinks that it is an imperfect duty to do so. He claims that, while it is possible to universalize 'Every man for himself' as a maxim,

> everyone who finds himself in need wishes to be helped by others. But if he lets his maxim of being unwilling to assist others in turn when they are in need become public, that is, makes this a universal permissive law, then everyone would likewise deny him assistance when he himself is in need, or at least would be authorised to deny it. Hence the maxim of self-interest would conflict with itself if it were made a universal law, that is, it is contrary to duty. Consequently the maxim of common interest, of beneficence towards those in need, is a universal duty of human beings.[25]

Yet Kant's argument refers specifically to people in need, and this criterion does seem to be important: while we might want to embrace the idea that there is a duty of beneficence in respect of need, the idea that we have a duty to give to those who are not in need, or are positively well-off, is much less intuitive, notwithstanding that to do so would be beneficent. So obligatory beneficence in this sense is perhaps better understood as being beneficence of a particular sort, directed to certain specific requirements. But if this is so, it comes up against my understanding of enhancement as being something that imbues attributes beyond what is recognized as being necessary for a reasonably flourishing life. In short, the Kantian argument might be able to establish that enhancement *qua* therapy is required; but about bare enhancement, it says nothing unless we think – implausibly – that everyone is in need simply by virtue of living a life that could conceivably be better in some respect. And while not every deontologist is a Kantian, Kant's account does seem still to be the best attempt to define a set of moral duties by means other than by *fiat*. More importantly, it is not at all clear that a general rule mandating the promotion of further welfare for the already adequately well-off would be forthcoming from any other quarter.

The remaining strategy would be a utilitarian one of welfare-maximization. On this account, moral demands are a matter of bringing about the best possible state of affairs in the world. Since beneficence makes at least some people better off, we ought to perform beneficent actions, at least for as long as it's efficient to do so and their benefit outweighs the cost of their performance. There is no particular reason why this policy has to be thought of as being restricted to the alleviation of need: since it is possible to improve the lot of the well-off, there might be a duty to do that. Yet it is not obvious that it really would be blameable not to improve the lot of the already-at-least-reasonably-well-off; and if it is not blameable, then not doing so would not be a duty.

Naturally, a reasonably sophisticated utilitarian argument ought to be able to distinguish between the obligatory things that we ought to do for the neediest, and the merely admirable things that we might do for those not in

specific need. The problem that this presents in respect of a putative duty to enhance is that avoiding the idea that we have the same obligation to benefit the non-needy as the needy requires drawing some sort of moral distinction between them; but since I understand 'enhancement' to refer to things that improve lives that would not otherwise be needy, it is not obvious that this sophisticated utilitarianism would have to, be willing to, or even be able to embrace enhancement as a duty.

## Beneficence and duties to enhance

There are other problems with the idea of the duty in a duty to enhance. Not the least of them is that, for quite different reasons from those suggested above, not all beneficent actions are obligatory, and their being hypothetically obligatory would not follow even from their being beneficent. So even if we do have a duty of beneficence, it does not follow that everything that makes the world a better place is a duty. This is because there is a number of ways in which a person may potentially make the world a better place, and potential duties may get in each other's way. Action *A* may make the world better than would *B*; in this case, there would be a clear moral reason to do *A*. But there may not be a *duty* to do *A*: it may be incompatible with doing *C*, which could make the world a better place to the same extent as *A*, although in a different manner.

Assume that *A* and *C* are unimprovable: there is no option *D* that is better even than they. In such a situation, we could say – based on an appeal to a duty of beneficence – that an agent would have an obligation to do either *A* or *C*. But we could not say that she had a duty to do *A and C*, since – although she has a good moral reason to do both – that would be impossible; and nor could we say which of *A* and *C* duty demands. (This provides a nice demonstration of the problem with Harris' understanding of the word 'duty'. If we accept that a duty is that which one ought to do, and that *ought* implies *can*, there is no obligation – no duty – to do *A* and *C*. But there is still a good moral reason to do both. Therefore what there is a good moral reason to do cannot be the same as what there is an obligation to do.) It would appear to follow that it is possible for an action to be world-improving but not obligatory. *A* and *C* are *potential* duties, but the mere fact that some action can improve the world will not mean that it is straightforwardly a duty. The most that we can to say is that, even if improving the world is a duty, a given strategy of doing so is mandatory if and only if it is the best possible way of improving the world to the same extent.

Does enhancement – of persons or of genomes – satisfy this demand? It is not obvious that it does. In fact, granted the many ways in which we could

improve the world in which we currently live, there is a good reason to believe that enhancement of human beings ought to make an appearance quite a way down our list of priorities, and that there is any number of ways in which not enhancing would not be blameable, so not violate any moral requirements. To give a couple of quick examples: it may be a very good thing to provide reading glasses to the presbyopic inhabitants of the slums of Lagos, but it is plausibly more urgent that we provide them with basic sanitation; it may be a very good thing to engineer a greater resistance to HIV in our descendant generations, but it is not obviously crazy to counterpose that with the thought that it is more urgent that we find a way to avoid global warming, develop safe nuclear fusion reactors, and so on. It is not obvious that providing or researching enhancements represents the most demanding or the most efficient of the world-improving things we could do – and it is a peculiar sort of duty that we invariably have a reason to abandon in favour of something else. (Of course, it may be possible to enhance in such a way as to dissolve many of today's problems. The absence of sanitation in a slum might be less serious were we to engineer immunity to cholera, for example. But this seems like a needlessly complex solution to the problem that would still leave millions of presently-existing people vulnerable, and I find it hard to accept the idea that we might have a duty to $\phi$ when $\phi$-ing achieves the same sort of thing as $\psi$-ing, but in a less effective and more complicated way; hence engineering cholera resistance is not plausibly a duty here.)

The most we can say is that enhancement is a *prima facie* duty – that is, the kind of activity that may be obligatory when there are no competing claims, but the blameability of which's absence may be nullified. Yet, pragmatically speaking, there will likely as not *always* be claims that not only compete for our attention, but are more pressing. The blameability of not enhancing will *invariably* be nullified. This means that, notwithstanding that enhancement would be a good thing, the 'duty' to enhance will invariably turn out to have been illusory.

The point is that it will not do to say that it is *always* blameable not to perform some particular action to improve the world: such a claim reflects, at most, only a particular manifestation of a more general, and implausibly general, duty of beneficence. If improving the lot of the needy – or the world as a whole, for that matter – can be better achieved by doing something other than *$\phi$-ing*, it may be permissible – and would probably be obligatory – not to $\phi$. And if this argument goes through, then it will apply to enhancement as much as to anything else. If a life that was not hitherto reasonably good can be made reasonably good by enhancements, then it might be blameable not to provide them – and so it might be a duty. But such enhancements do not meet the everyday meaning of the word 'enhancement' to begin with. It's not nearly so clear that anything more would be blameable.

## What is enhancement?

So there's a problem with the claim that enhancement is a duty that stems from how we're supposed to understand 'duty'. If there's ever a situation in which *not* to enhance would be at least as acceptable a course of action as enhancing – and I posit that there frequently, if not invariably, would be – then it would not be obvious that enhancement is a duty at all.

It turns out that the word 'enhancement' is at least as problematic.

To begin with, the word needs fleshing out. As I hinted a little while ago, to say that enhancement is a good thing is to say very little, since if an alteration – whether to a gamete, an embryo, or an adult – turned out not to be a good thing, we would thereby lose the entitlement to call it an enhancement. At the most, we would say that we had *attempted* to enhance this or that characteristic, but that we had failed. (Note that one could *attempt* to enhance without ever once actually managing to improve the world or even much expecting success. But there being a duty to *attempt* enhancement is not the same as a duty to enhance, and we can't say that it's attempts to enhance that are obligatory without relinquishing our claim that it's welfare that provides the moral motor of obligation.)

Harris' idea about 'enhancement' meaning anything that makes a change for the better does improve things – but it is still rather formalistic, since it leaves the question of what counts as an enhancement wide open. For example, someone could quite understandably include on the list of enhancements things like increased intelligence or athletic ability. Equally, heightened resistance to malaria or the common cold, or a reduced chance of suffering from irritable bowel syndrome, would also be fairly straightforwardly classifiable as enhancements. However, giving someone the ability to echolocate would be to make a change and, assuming that nothing comparable was lost in the process, to make a change for the better – it would make navigation around the house at night a lot easier – and so engineering that into humans would plausibly count as enhancement. The same might be said for any number of other characteristics, too.

Here lies the problem: there is an indefinitely long list of things that might count as an enhancement. But, if there is an indefinitely long list, then it is hard to sustain the idea that providing those enhancements might be a matter of duty. After all, if the list of enhancements we could provide (and are in some sense obliged to provide) is indefinitely long, then it means that, if we do nothing, there will be an indefinite number of things we ought to do. But this will be no less true if we do *anything*, or even *everything possible*: there will still be an indefinite number of things we ought to do. In other words, no matter what we do, we will not be any closer to discharging our moral duty – and this makes it hard to see why we must even try. It would not be reasonable to blame someone for not having met every demand on an indefinitely long list – and so a duty that is impossible to discharge doesn't look like much of a duty at all. Call this the *Endless List Objection*, or 'ELO' for short.

There is a few counterarguments to the ELO. First, and apparently strongest, is to point out that a duty to enhance is not the same as a duty to provide every possible enhancement, any more than a duty to rescue a drowning child is a duty to rescue her in every possible way. There being lots of ways to rescue the drowning child doesn't mean that we can avoid choosing one of them; that there's a lot of ways to enhance doesn't alter the obligation to enhance. But this counterargument doesn't quite work – not least because, once the child is rescued, the duty definitely *is* discharged and it is no longer incumbent on us to do anything. This is not true in respect of enhancement; here, we are in the position of the Red Queen, for whom 'it takes all the running you can do to keep in the same place'.[26] A cognitively enhanced human could still be athletically enhanced, kept young, immunologically enhanced, cognitively enhanced a bit more, given the ability to echolocate, and given another top-up of cognitive enhancement. The point is that it is hard to see how the notional duty to enhance could *ever* be discharged. The list stays endless.

This leads on to the second counterargument, though, which points out that *even if* the list of possible further enhancements is endless, there comes a point at which it is reasonable to stop providing them (just as it is reasonable to stop attempting to rescue the child after a certain point). But while this is true, it gives the game away, because we are no longer talking about a duty to enhance *simpliciter*: we're now opening the way to a set of claims about what is reasonable enhancement. I'm inclined to think that all enhancement is more than reasonably demanded; but even if that claim falters, there would still be space to talk about how many enhancements are enough. So while it's true that one would be able to avoid the ELO if the argument was about a more nuanced duty to provide enhancements, making this move would simply raise another requirement that the proponent of the position be willing to argue that *this or that* specific enhancement is a moral duty.

Aside from the ELO, an interesting thing to note about Harris' claim in particular is that it is not that we should enhance human *characteristics*, but that we should enhance *humans*. This point raises interesting questions about the nature of the enhancements that we're supposed to instil.

Harris cannot say that what he means is the enhancement of human characteristics without undermining his own argument fatally, since humans are characteristically intelligent within a certain range of intelligence, characteristically athletic within a certain range of athleticism, characteristically susceptible to malaria within a certain range of susceptibility, and – importantly – characteristically subject to the whim of fortune when it comes to stepping outside these ranges. If characteristically human traits are the mark by which we are to understand what an enhancement is, there is no need for us to enhance evolution: natural selection has already made us as characteristically human as we ever could be.

Nor could he appeal to non-characteristically human traits to provide the mark of what counts as enhancement. After all, although we could strive to ensure that each child in some future generation combine the various virtues of the brightest and best of humanity as it stands at the moment, even this seems to be aiming low. Consider, for example, intellectual enhancement. It is possible that no human – even uncharacteristically clever humans – will ever be clever enough to reconcile quantum with relativistic physics; it is possible that one would have to have a different sort of brain altogether.[27] Moreover, there's no reason to think that such intellectual feats would represent the epitome of cleverness anyway. So if we're going to enhance at all, it is not clear why the abilities of even the most extraordinary members of the current crop of humans should represent the limits of our ambition: we would seem obliged, by the lights of at least some enhancement enthusiasts, to pursue them beyond even abnormal human ranges. (This is why Harris seems to be, on the whole, a trans-humanist.[28] I say 'on the whole' because Harris is not always wholly clear on the matter: in Chapter 1 of *Enhancing Evolution*, for example, he talks about enhancements that will help humanity 'achieve its potential'[29] – though, in the light of the considerations I have just presented, that seems rather paltry.)

In other words, when it comes to enhancement, the reason to provide non-characteristically human traits seems to be displaced by the reason to provide characteristically non-human traits. But this takes us straight back to the ELO: we could always be a little bit cleverer, faster, better-looking, or more creative; and if there actually was an upper limit on the number of enhancements that we could provide within the bounds of humanity, there certainly isn't if we've decided to break through them. Certainly, the onus seems to be on Harris and those who share his brand of genomic optimism either to bite the bullet and admit that the scope of the enhancements that we are ostensibly duty-bound to provide *is* pretty much limitless, or else to dodge it by providing some criterion of genomic propriety that does not undermine the programme of enhancement teleologically before it ever gets going – and that looks difficult. After all: if we are willing to say that there is some sort of upper limit on the amount of enhancement that we are obliged to pursue, we would have to be able to say what it is without opening the possibility that we have, as a species, reached it already.

The point is that, within a phrase like 'enhancement is a duty', the word 'enhancement' is at least as problematic as the word 'duty', and we should therefore be wary of it. One might wonder what remains.

## What is 'acceptable'?

The argument of the last few sections would apply to anyone who claims that there is a duty to enhance. Harris' claims provide a handy lens through which

we can examine the structure of a kind of argument that is not uniquely his; but any person occupying a comparable position would have to be able to give a compelling definition of duty, and be able to give an account of how there could be a set of bare enhancements that are obligatory, but which is constrained tightly enough as not to fall foul of the ELO nor so tightly constrained as to admit that the only obligatory enhancements are those that are justified by their ability to make a less-than-reasonably-good life good enough – and so fall squarely within the category of therapy.

However, there is a few things that Harris says that do seem to be more specifically *his* claims, and which merit attention; I'll try to widen the scope of their relevance here, but cannot claim that my arguments over the next few paragraphs will have a guaranteed wide application.

Harris seems in some places to have fallen victim to linguistic ambiguity. A nice example of this is provided in the context of his discussion of the acceptability of death from old age. Consider again what Harris has to say on the matter:

> Dying of old age is acceptable only because it is regarded as inevitable. Absent the inevitability, the acceptability is problematic.[30]

I'm going to look more closely at senescence, and the possibility that we might make use of the biosciences to eliminate old age, in the next chapter; it's the word 'acceptable' that's of concern here. This can bear being read as a component of an argument for death from old age (or for any number of other bad things) being unacceptable.

Does the argument pass muster? I think not. The argument is weak, and its weakness is caused by an equivocation around the word 'acceptable', which could mean either 'that which is capable of being accepted' or 'that which is worthy of acceptance'. A student's essay may be acceptable in the sense of it being possible for me to accept it, but unacceptable in the sense that it isn't very good. The two senses are very different; it is the latter sense that tends to dominate in everyday speech, and Harris admits that it is in that normative sense that he is using the word.[31] But this normative sense does not belong in any claim about dying of old age being acceptable because inevitable. Of course, from the fact that people *do* accept age-related death, it must follow that they *can* do so; and, from this, it is 'acceptable' in the (non-standard) descriptive sense of the word: it is capable of being accepted. But it would be no less acceptable in this sense if death from old age was something we could prevent easily: a person's accepting such a death is not impossible. Even if no one *would* or even *ought* to accept it, they *could*, and the argument fails as a result.

If the argument is to make sense, we have to read the claim in the sense that dying of old age is not worthy of acceptance. But why on Earth we should believe that death from old age is unacceptable in the normative sense? To say

that something is unworthy of acceptance in this way is to appeal to a norm of some sort that has been violated: the unacceptable essay has violated some rule on plagiarism, the unacceptable behaviour some rule of social propriety, and so forth. But what rule has been violated here, and by whom? If an essay is unacceptable, it is the author who has failed to satisfy some standard. But, when someone becomes senescent or dies from old age, the same does not apply: we aren't looking the Reaper in the eye and telling him he has broken the rules. And if someone said that he was willing to accept the possibility of dying of 'natural' old age 50 years before technologically necessary, we would not – presumably – want to say that *he* had broken them, either. His death would be normatively acceptable to him, and it would be at the very least odd to think that he was mistaken on the matter.

Is it, then, the rule that tells us we have to save lives if and when possible that means that it is unacceptable for us to stand and watch as someone places his first foot in the grave? (We do say that to stand and watch while a child drowns is unacceptable: mightn't the same apply here?) Assuredly not. That there's a disanalogy between preventing death from old age and preventing death from drowning is something I'll demonstrate in the next chapter. But even without that, the claim would have to apply either to people who currently exist, or to potential people, or to both. To advocate enhancement on behalf of those who already exist as a means of saving their lives is either doomed to fail (since, as we have seen, granted that we ought to save lives if and when we can, it does not follow that we have a duty to enhance when there are more effective ways to save possibly more lives) or strange (either because it is hard to see why we are obliged to intervene to save someone's life if that person is perfectly happy to accept his death, or because it suggests – implausibly – that pulling drowning children from lakes is enhancing them in some way). So the rule must apply, if it applies at all, to those people who have yet to exist.

But this also cannot be. The Harrissian argument at this point would rely on a form of the PPB to provide its motor. But who benefits from preemptive enhancement? Noone. If we manipulate the gametes that will become an embryo, we are not improving the life of anyone who currently exists so much as changing the range of possible people they may become, and so providing some possible person with his only chance of becoming an actual person. The same applies if we are trying to choose between implanting an enhanced and a 'natural' embryo after IVF, since implantation is an embryo's only shot at having the kind of life that could be good or bad in the first place. If my only chance at existing was in an un-enhanced state – as it was – then I would *not* be benefited by the cells-that-became-me having been manipulated, because, had they been manipulated, *I* would never have existed at all. Someone similar to me in many ways would have, and that is all very well for him – but he would not have been me. I have not benefited from enhancement, because I only ever had one shot at existence anyway – something that Harris admits cheerfully.[32]

By contrast, things are not the same when we pull the drowning child from the pond to save her life. Since she does already exist, our decision is the grave one of whether or not she, a creature who already exists, should continue existing. The decision concerning whether she should *begin* to exist in the first place is not raised. Suffering a person to exist who is vulnerable to age is simply not the same as refusing to pull a child from a pond; it is more like suffering a person to exist notwithstanding that they might drown. One would not normally want to label the possibility of such an existence unacceptable.

The conclusion that seems unavoidable is that the claim about age-related death being unacceptable must either be false or meaningless, there being no violation of any norms. *Acceptability* provides the wrong scale by which to measure senescent death. Such a death is not like the essay that we don't accept because it is hand-written on the back of an envelope. And so this aspect of Harris' claims about enhancement being morally demanded is not compelling.

## A duty to enhance?

Ensuring the best life for themselves and their offspring is a project in which humans have been engaged, wittingly or not, for millennia. Indeed, any animal that prefers to mate with the strongest and healthiest could be said to be responding to a genetic imperative to seek out the fittest for the sake of the fitness of the next generation; a selfish gene 'wants' to ensure it has the best vector possible. And since enhancements contribute to this project, and are by definition good, it seems fairly straightforward to say that the project of enhancement is good as well. However, for that project to be a *duty*, it would have to be more than a good thing: it would have to be shown that agents are obliged to do positive good such that they could be legitimately blamed for not doing so, that enhancement is the *best* way to do this – lest there be times when we had a duty *not* to enhance, but to do something else – and that it would be possible to discharge the putative duty.

I have argued that, on all counts, arguments about there being a duty to enhance fail. Reading 'enhance' in broad terms means conceding that there is an indefinite number of potential enhancements we can make. No matter what we do, there will always be something else that we could yet do in addition. And this, I have claimed, looks odd.

We could try to interpret the duty narrowly, and talk about specific enhancements that we are obliged to provide, or specific facets of the project of enhancement that we are obliged to pursue. But now – bearing in mind that this is already a large move away from the initial claim that enhancement in the wide sense of the word is a duty – we have to decide *which* enhancements demand our attention, and in what circumstances. My claim is that, since to be un-enhanced is not the same as being in a harmed state, if any kind

of enhancement did turn out to be mandatory, it would be the kind that is described in ordinary language as therapeutic – which is as much as to say that it would not be much of an enhancement as the word is commonly understood at all. But even ignoring that claim, it seems not to be improper to suspect that at least many things are more demanding of our attention than at least most enhancements. Even allowing that there is a duty to improve the world, it remains to be shown that any particular enhancement – let alone enhancement *simpliciter* – would be the best possible way of doing so. This is a very tall order at best; in all likelihood, it is impossibly tall.

Enhancements are good – that much is analytically true – and they might very well increase a person's flourishing. Providing them is a good thing. In fact, it seems plausible to suggest that providing enhancements is more than good – it is admirable. But it is not a duty. In saying that it is, one would be overplaying a hand that was plentifully strong to begin with.

## Notes

**1** Kant, of course, thinks that there are duties to self (Kant 1998, 6:417 ff); but a duty to self-enhance would not be among them – the closest he'd get to that claim is the claim that there is a duty to develop one's talents (Kant 1993, 4:422–3). See Timmermann (2006) for elaboration.
**2** One interjection here might be that vaccinations can be thought of as enhancements, and we might have a moral duty to vaccinate ourselves for the sake of public health. If this is permissible, then the back of the totalitarianism objection has been broken. However, there does seem to be a response to this, in that a policy of requiring vaccination has a couple of aspects. In the first, it is an enhancement for the sake of the person being vaccinated – and none of what I say in the coming pages makes this problematic. In the second, from the point of view of the public, vaccination perhaps isn't so much an enhancement as a means of minimising the risks faced by others. Whatever our opinions on such a policy, they won't tell us anything about enhancement.
**3** Savulescu, 2001, *passim*.
**4** Savulescu and Kahane, 2009, *passim*.
**5** Bennett, 2009.
**6** Herissone-Kelly, 2012, *passim*.
**7** Vide Herissone-Kelly, 2006, *passim*.
**8** Harris, 2007, p. 36.
**9** Ibid., p. 29.
**10** Ibid.
**11** Ibid., p. 29; cf. pp. 14–15.
**12** Jensen, 2011, *passim*.
**13** Ibid., p. 39; emphasis mine.
**14** Ibid., p. 52.

**15** Ibid., p. 53.
**16** Even in cases of voluntary euthanasia, I take it that death is not desired on its own terms. The best we can say about dying is that it is better (by some standard) than the alternatives in some cases – but that is hardly sufficient to demonstrate that mortality is worthwhile.
**17** Ibid.
**18** cf. Harris, 2005, pp. 242–8; Harris, 2007, § 11.
**19** Personal communication, 4 March 2008.
**20** Ross, 2002, *passim*.
**21** Beauchamp and Childress, 2009, *passim*.
**22** Harris, 2007, p. 50.
**23** Ibid., p. 189.
**24** Harris, 2007, p. 3.
**25** Kant, 1998, 6:453.
**26** L. Carroll, 1998. *Alice's Adventures in Wonderland/Through the Looking-Glass*. Oxford: Oxford World's Classics, p. 145.
**27** If, astonishingly, this problem has been solved by the time this is published, substitute some other conundrum such as $P = NP?$.
**28** Ibid., pp. 39–40.
**29** Ibid., p. 11.
**30** Ibid., p. 52.
**31** Ibid., p. 53.
**32** see Bennett, 2009, *passim* for an expanded critique of the principle of procreative beneficence; cf. Harris, 2001, p. 385 and Harris, 2007, pp. 105–6.

# 4

# Sex, Death and Cabbages: A Defence of Mortality

Having argued that there is no duty to enhance others, it might still be something that it is a good idea to do; and it might be a good idea for us to enhance ourselves in various ways, too. Over the course of this and the next three chapters, I'll examine this idea, and defend the claim that, while most potential enhancements are permissible, they are also much less likely to improve our lives than we might hope. A small category of enhancements *does* stand a reasonable chance of promoting flourishing; counterintuitively, these are enhancements that look to be at first glance the most trivial. And there is one kind of enhancement that probably ought to be more positively resisted, as I'll argue in Chapter 7.

In this chapter, I'll take as my test case the possibility of significant life extension – which is something that I take to be the *excelsior* of enhancements that we might want to make to the body. John Harris claims that 'life-extending therapies are, and must always be, lifesaving therapies and must share whatever priority lifesaving has in our morality and in our social values'.[1] For him, there is no difference between saving someone's life, and extending that life, because saving a life just is a matter of extending it, and – '[s]ince lifesaving is just death-postponing with a positive spin'[2] – *vice versa*. If we think that saving lives by curing illness is a good thing, that's all we need to be committed to the idea that significant life extension is a good thing. Aubrey de Grey agrees: for him, to show that there is no meaningful difference between saving and extending a life is easily done,

> because it derives directly from the definitions of those concepts. In each circumstance we are giving the beneficiary a greater remaining healthy potential lifespan than they would have if we held back, which is the beginning and end of what we mean when we say we have saved their lives, and also of what we mean when we say we have extended their lives. This is not a matter of the absence of a contrary argument; it is a matter of definition, plain and simple.[3]

And – to be fair – there does seem to be something commonsensical about the position. When we save the child's life by pulling her out of the pool, we are extending it; and the desire that she not die – that dying is bad, and that it should be averted – is presumably an important factor in our having decided to save her life. Ask many people what they think the first virtue of medicine

and medical research is, and they'll likely as not tell you that it's to save lives – which is, it would appear, as much as to say that it's to avoid deaths. Indeed, the NHS distributes resources in a manner that pays attention to the number of life-years, suitably adjusted for quality, that they'll generate. Harris, as we've seen, thinks that we accept dying of old age only because we see it as inevitable;[4] de Grey, in only a slightly more radical claim accuses many of us of being in a 'pro-aging trance':

> [A]ll the major current reasons given for not curing ageing are mere crutches to help us cope with the immutability of ageing that we have been brought up to accept. Once we realise this, our determination to consign human ageing to history will be second only to our shame that we took so long to break out of our collective trance.[5]

Nevertheless, the point of this chapter is to argue that, while protection against ageing is probably a desirable thing, significant life extension as an end in itself probably isn't. It might even be positively undesirable. Moreover – and this is, I admit, a slightly more hesitant claim – once we have got rid of the error of desiring to eliminate death, the problems presented by old age can be dealt with in ways that do not rest on the idea that it is incompatible with a good life.

## Defending against death

It's fairly obvious that, if we're concerned with promoting human flourishing, that the humans in question should be alive is an important consideration. So it makes sense that significant life extension should figure highly on our list of potentially important enhancements. Being dead may not be bad for the dead – and I'm going to accept that Stoic idea as a given for this chapter (death, says Seneca, 'would have been worth fearing [only] if it could coexist with you'[6]) – but the point of death is also fairly clearly the point at which any kind of flourishing is terminated. Moreover, as Harris points out, '[g]iven that people want life and fear death it is difficult not to see longevity, and perhaps immortality, as a palpable good'.[7] Woody Allen, for his part, isn't scared of dying; he just doesn't want to be there when it happens.

And we are (at least if people like de Grey are to be believed) tantalizingly close to living in a world in which natural deaths could be eliminated. Though there are some diseases that we cannot cure at present, we can be fairly confident that they are not incurable or ineliminable in principle. HIV presents us with a major source of morbidity, but it is, in the end, only a virus. We have already made great progress in turning HIV from something that kills with reasonable certainty and in a reasonably unpleasant way, into a chronic-but-manageable condition that need not make all that much difference to a

person's quality of life or death; and the more research is done, the better our chances of finding a cure or vaccine.[8] The same applies to other diseases that, though currently incurable or unpreventable, are neither incurable nor unpreventable in principle. In all these cases, a great deal of suffering can be avoided; but not only that: a lot of deaths can be avoided, too.

The natural ageing process may strike us as being a little different; ageing is a part of what bodies do, and is therefore not so easily presented as being a result of an invasive organism or a cell gone haywire. Nevertheless, the process of becoming elderly is, when it comes down to it, just another biological process; and, as such, it ought to be possible to interfere with it. To some extent, we do that already: something like HRT is designed to interrupt an aspect of natural ageing; other drugs are designed to maintain heart function, and so on. The next step along this path is the use of bioscientific insight to repair some of the damage caused to the body by the ageing process 'from within' – for example, we might imagine the use of stem-cell therapies to rejuvenate flagging heart tissues. Granted the fact that nature manages to create tissues and organs anew from a basic genetic recipe, the general idea is that there is no reason in principle why we shouldn't be able to do something similar for ourselves, and either provide ourselves with endless spare parts, or – better yet – work out how to make ourselves less likely to wear out in the first place.

While there is a whiff of the science-fictional about many of these procedures, they are all things that have been suggested seriously as potential areas of research and as procedures that we could see beginning to be realized within the tolerably near future. It is at least conceivable that we could make use of the biosciences to make ourselves, to all intents and purposes, immortal. 'Immortals' in this sense would not be immune to dying – people whose bodies had been primed against the ageing process would not, of course, be immune to road accidents – but they would be resistant to dying of old age, and perhaps more resistant to dying of some 'non-natural' deaths, too: a road accident that causes a 20-year-old little more than a few bruises might be a causal factor in the death of a 90-year-old today; but if her body could be made less vulnerable generally, she would be in much less danger from it.

Added to this is the fact that the longer someone lives, the greater the chance that they'll survive to see some medical innovation that can help them live longer still. Aubrey de Grey has claimed that the world's first 1000-year-old could well be not only alive, but an adult, today; part of the reason for this is that, though we do not have the ability to make someone live that long right now, we do have the ability to keep anyone who is less than about 60 or 70 years old alive, and in reasonably good health, for another 20 years – and there is nothing difficult about that. With the right kind of research input, by the end of that 20-year period, we could easily have the wherewithal to maintain their life and health for a much longer period again. The process would be something like the reverse of Zeno's arrow: whereas the gap between the arrow's head

and its target narrows forever without closing despite the arrow's constant movement, so the gap between the currently-60-year-old Immortal *in embryo* and his death widens forever despite his. Ray Kurzweil and Terry Grossman, after having claimed (somewhat perplexingly) that we are built according to 'an outdated genetic program that wants us to die young'[9] describe affairs thus:

> [W]e still don't possess the full technology an skills to enable us to [overcome our genetic limitations], but it is very likely that we will have them in the near future. Therefore, our goal is to *live long enough* (and remain healthy long enough) to take full advantage of the biotech and nanotech advances that have already begun and will be occurring at an accelerating pace during the next few decades. If you stay on the cutting edge of our rapidly expanding knowledge, you can indeed *live long enough to live forever*.[10]

Assuming that Immortals would have survived the killers of infancy, and would not have to face the killers of old age, we can do a little actuarial work to determine the death rate for average adults, and use that information to make predictions about the attrition rate among Immortals. Hence John Harris quotes Steven Austad as having suggested that the average life-expectancy for Immortals could well be just a little short of 5000 years[11]; that is, some would be struck by lightning on their 30th birthday, and others would live much longer – but it would not be unreasonable to expect to go around 5000 years before experiencing a fatal misfortune.

Is it really better that we defer death, or do not die at all, though? That's the question that I'll address in this chapter. I'll defend the position that Immortality would not be something that a reasonable person would list as a characteristic of the good life.

## Avoiding deaths and saving lives

One point needs to be got out of the way before engaging with the question in earnest, though, and it speaks to the therapy/enhancement distinction. As I conceded above, medical science has become very good at saving our lives; as such, it has done its bit to make us live longer. An Immortal is, by definition, someone who lives a great deal longer.[12] But beyond this, there is nothing all that special about an Immortal; and it might well be assumed that a great deal of fairly workaday medical treatment nudges us some way in the direction of becoming Immortal. In turn, this might be taken as an indication that anyone who supports the progress made by the biosciences in increasing lifespan (and health span) ought, on pain of inconsistency, to support living longer (with a correspondingly high quality of life) outright – and, by the same token, that to express doubts about the Immortalist programme is to express doubts about medicine more generally. But this is not quite correct.

A great deal hinges on providing an account of what is accepted as a reasonably long life. As Thomas Nagel has observed, we do tend to think that Keats' death at 24 years of age was somehow a greater loss than Tolstoy's at 82.[13] Having said this, as he goes on to point out, though dying at 82 is seen as less bad than dying at 24, it might still be terrible to die at 82, or 806. Nevertheless, the idea that the badness of dying diminishes the longer the life in question is the kind of sentiment that is traded on by the so-called 'fair innings' argument, which suggests that – at least when it comes to the allocation of scarce health resources – the young trump the old. This only makes sense if we think that there is a point beyond which people can be said to have had at least a minimally acceptable lifespan. But thinking about minimally acceptable lifespans brings us back to what makes the difference between 'therapeutic' life-extension and Immortality: the reasonable expectation criterion outlined in Chapter 1. If one accepts that there is a reasonable expectation concerning survival times, then this is all one needs to be able to draw a line between interventions that aim to ensure that people live to a minimally acceptable age, and the kind that aims to ensure that they live a great while longer. Suspicion of radical life extension does not imply suspicion of medicine, even if we think that the point of medicine is to extend life.

Does this mean, then, that it would be acceptable to leave a centenarian to drown in the pond on the grounds that he has had a reasonably long life and then some? Actually, it doesn't; and this is because there is a difference between saving someone's life by rescuing them, and postponing death in the sense that we are considering here. In the normal run of things, the idea that one might save a life by means of rescue is reserved for those situations in which a person is faced with some imminent threat. This is the situation in which the drowning centenarian finds himself. By contrast, one does not rescue anyone by providing swimming lessons, or by vetoing the construction of the pool to begin with. (This is a point that will have a role to play in Chapter 9, too.)

Now, I understand terms like 'significant life extension' to indicate procedures that aim to defer the point of death for as long as possible by minimizing certain threats – if possible, by making sure that they never arise. This is not the same as rescue. One might plan to make use of a certain life-extending intervention next Saturday and book an appointment to see the relevant expert; it's strange to talk about a person planning and making an appointment to be rescued then. It *might* be that the apparent difference between the two ways of describing the appointments is purely illusory; but this is not something that we can take for granted – and the two descriptions do not appear to be mutually substitutable. On these grounds, and notwithstanding that both life-saving by rescue and by significant life extension refer to reducing death tolls, it remains possible coherently to stick to a claim that there is an obligation to attempt the rescue of centenarians in danger – which does, of course, amount to an attempt to avoid their dying – without having to admit that whatever we say about cheating death in this situation applies *pari passu* to cheating death in all others.

Moreover, questions about the morality of rescue are, in a certain sense, differently focused than are questions about extending lives. The ethics of significant life extension revolve largely around the person whose life is potentially going to be lengthened – around what *he* wants, what is reasonable *for him*, and so on. By contrast, asking these questions in respect of someone whom we can rescue is an indication that we have had, to borrow Williams' phrase, several thoughts too many. Rescuing someone has less to do with the details of their life and assessments of what they want and what is reasonable for them, as with the details of ours – our relationship with the person rescued, perhaps, or the preservation of our own sense of basic decency and moral character.

It might be interjected that there are all kinds of reasons to prefer never to be in a situation in which someone requires rescue. If we had the means to rejuvenate the muscles in the heart, for example, we might thereby save people from having to be rescued from heart attacks. It might be that this is, in the long run, more desirable. But that doesn't alter the point that there is a coherent distinction to be drawn between rescue in the everyday sense of the term, and the indefinite avoidance of mortality. One cannot leap directly from the former to the latter. Avoiding death will lead to longer lives; but taking pains to prevent a death is not the same as promoting Immortality because there is a class of deaths that one might try to prevent without being engaged in the project of indefinitely avoiding death.

## What's wrong with mortality(?)

If there is a coherent difference between having one's life extended by means of rescue, and having it extended by means of coming never to require rescue, it means that a person may endorse being saved from imminent danger, and may endorse Immortality, but that endorsing one does not imply endorsing the other (although someone who endorses immortality for himself is presumably going to be *de facto* the sort of person who would want rescuing as long as there was a reasonably high quality of life assured for him afterwards). Would a failure to endorse indefinite life extension be wildly eccentric, though?

One way to begin to answer this question is to look at why dying is taken to be bad. We can account for the badness of dying by saying that it deprives us of the goods of life. This is a position that has been occupied by several people. Death represents, for Nagel, the 'abrupt cancellation of indefinitely extensive possible goods'[14]; Williams articulates the position by saying that

> to want something is to that extent to have reason for resisting what excludes having that thing: an death certainly does that, for a very large range of things that one wants. If that is right, then for any of those things, wanting something

> itself gives one a reason for avoiding death. From the perspective of the wanting agent it is rational to aim for states of affairs in which his want is satisfied, and hence to regard death as something to be avoided.[15]

Death deprives us of things we like and want; therefore it makes sense to consider death to be bad and undesirable; therefore it would seem to make sense to consider the opposite of death – extended life – to be good and desirable. Of course, a great deal hangs on the goodness of the thing lost: if, for some reason, we no longer regard our lives as good, then death may not be as unwelcome as it could be. But the basic message is intuitively powerful: it is better for us not to die, or at least to defer our deaths for as long as possible, because life is a necessary condition for us doing things that are important to us, and its termination means losing the opportunity to do them. Let's call this the 'Lost Opportunity Thesis', or 'LOT' for short.

As it happens, I think that the LOT is about as good an argument for Immortality as one is likely to find. Its power is significantly related to the idea that, for a life to be desirable, it must be *for* something, and that something has to be more than yet more life. If we imagine a person whose only project is continued living, and ask him why he wants to stay alive, and he tells us that it is because that way he will be alive longer, there would be some justice in our supposing that he had not quite understood the question: a project cannot easily be accepted as its own justification. A longer life will not provide a reason to live longer; it is hard to see how the phrase 'At least you're alive' would be of much compensation to someone who was deprived of all *but* his life. I am, for this reason, satisfied that life for life's sake is – certainly from the point of view of the working hypothesis that flourishing is important – a non-starter.

This point is, I think, almost a platitude. It can be deepened a little by pointing out that what we *do* value about being alive may not always be compatible with the strategies that we may have to adopt in order to achieve more of it. Hence it might not only be the case that a life lived for the sake of extending life is at best dubiously worth living: it might easily turn out that that life would be less desirable than a life not so directed. For example, one of the strategies that have been suggested as contributing to an extended life is caloric restriction. And, more generally, it may well be possible to adopt all manner of modified behaviours that will help us to live longer. For example, in a *New Scientist* article in 2010, Ray Kurzweil claimed to have all-but-cured his own diabetes by coming up with his own, strict, nutritional regime – a regime that includes consuming only 1,500 calories per day, based on under 80 grams of carbohydrates, and 150 supplements daily – all topped up with weekly intravenous infusions.[16] Kurzweil and Grossman, in *Transcend*, recommend avoiding many kinds of food for the sake of a longer life.[17]

Let's allow that these strategies for staying alive actually are effective. The problem with this is that eating is one of life's pleasures (and it is not for nothing

that deriving all our nutritional needs from pills is an idea that never caught on); a rich, varied and plentiful diet is *fun*. Hence it is no small question to wonder whether sacrificing that pleasure upon the altar of extra life-years is worth it. In considering this question, we shouldn't forget that the sacrifice would not be a one-off event followed by many extra years of cakes and ale, either: we would be talking about extra years of abstemiousness. The strictness of the Kurzweil regime is no small matter. Kurzweil and Grossman even suggest that

> tap and bottled water often contain environmental toxins, so we believe it is important that you filter tap water and know the quality and source of bottled water. [. . . I]t is important that you filter the water from your tap before you drink it.[18]

Alternatively, you could drink when you're thirsty, and not worry about it. Even if tap water is not as pure as it *could* be, it is not unsafe – at least, not in the parts of the world where books like *Transcend* are likely to sell well – and so that notion that it is *important* to filter tap water does seem to be a little excitable.

Assuming that the science behind this kind of suggestion is sound, adopting the kind of lifestyle required to reap the rewards is rather burdensome – and though these burdens may be adopted into a daily routine that would become just as natural as does anyone else's if all else were equal, all else is not equal, because the world is (generally) not geared to that kind of regime. Moreover, too close an adherence to this kind of lifestyle also risks becoming antagonistic to flourishing because of other aspects of the psychological burden that it implies. One characteristic that we might reasonably expect to be found in a flourishing life, for example, is some capacity for spontaneous action – action with which rigid adherence to any kind of regimen is obviously incompatible. It is not at all evident that life for its own sake is worth making this kind of sacrifice, or that an extended period of this kind of self-discipline is all that attractive an option. And adapting the regimen to be compatible with an enjoyable life means that rules become guidelines, and guidelines become good intentions, and good intentions get measured against other good things, and we find ourselves living the life we had before.

So: I am satisfied that, if Immortality is desirable, it is so because of some version of the LOT. This is not enough to show that it actually is desirable, though. There's a number of arguments that might purport to show that it isn't.

## Why not be immortal?

The first kind of reason that might be produced to resist significant life extension has to do with worries about what kind of world it would be that had Immortals in it; the second has to do with the kind of life that would be lived.

## The overpopulation objection

One of the arguments against the pursuit of immortality is that it would lead to a massively overpopulated world; a world in which people do not die is one in which more people are forced to live side-by-side, simply because nature has not managed to eliminate the bulk of them. This means that there will be many more people competing over the same resources, and polluting the world in the process, and a corresponding diminution of the quality of life for each. This is not in itself sufficient to show that there is anything wrong with pursuing significant life extension, for as long as the life of Immortals is just about worth living.[19] All the same, a life that is just about worth living is not really a *good* life, and it's the possibility of postponing death indefinitely for the sake of a good life that concerns me here – so some version of the argument does seem to stand.

Yet the argument can be countered. Importantly, 'overpopulation' is a term that really only makes sense within a context: there is no fixed quota of humans that the planet can support. If the planet's carrying capacity can be increased, there ought to be little reason to worry too much about overpopulation. And it is far from a given that the planet's carrying capacity cannot be increased. The problems of guaranteeing a sufficient supply of food, for example, are not insurmountable; even now, there is a great deal of food that goes to waste already that could be used to feed the hungry, and many of the problems of feeding an increasing and increased population can be solved by means of wise economic management. A much more powerful point could be made alongside this one, which is that a population that has the means and the will to carry out the kind of cellular and genetic tinkering necessary to make people Immortals presumably also has the wherewithal to tinker with the means of agricultural production. Problems of undersupply could, at least in principle, be solved by a new agricultural revolution that puts effort into the development of, and commits to the use of, crops capable of producing food from what is now marginal or unusable land. It might be necessary, alongside this trend, for people to become vegetarian, or to adapt to a diet based on, say, mealworms – getting protein from meat in the familiar way is highly inefficient – but the point stands that, with the application of the kind of scientific know-how necessary to create Immortals in the first place, it would be possible to create the food resources necessary to feed them. And much the same could be said in respect of other resources: a world with a higher population might quickly become intolerable if we rely on the same energy economy; but this is a long way from meaning that a radically different energy profile, that develops and makes much more use of renewables, and novel fuel sources such as biodiesel from GM algae, or even nuclear fusion, would not be possible.

Moreover, a world with more Immortals might turn out to be better in important respects than the world in which we live now, because Immortals

would have an incentive to take care of the planet that we do not. Though it might be unjust for me to act in a way that will make life harder for my great-grandchildren – say, by refusing to minimize my $CO_2$ emissions – there is little reason apart for the call of my conscience to do anything about it, since I will be dead by the time the consequences of my behaviour manifest themselves. Mortality gives me an escape route by which I can avoid behaviour that we might think is required. Immortals, though, would have a much greater incentive not to pollute or overburden the planet in other ways, since someone who can reasonably expect still to be here in 800 years' time can, by the same token, reasonably expect to have to carry the burden of his actions.

Another guard against overpopulation could be that we impose a *quid pro quo* – that those who would undergo whatever the procedures are for Immortality should be required to reduce their fertility. If we assume that Immortals would simply be people whose ageing process had been halted, they might well be expected to be fertile not just for decades, but for centuries. This would mean that they could be expected to have many more children over the course of their lives. (Many may prefer to have the same number as they would have had in a normal lifespan; but there would be at least some who would continue to procreate, if only because the chance of becoming accidentally pregnant during your life is much higher if the fertile part of your life lasts the best part of a millennium, and some of those accidental pregnancies would doubtless be carried to term.) Moreover, a rise in the average number of children born to each woman would be compounded by the fact that previous generations would not have died off – so the population would still be rising. Even if we can increase the world's carrying capacity, there may still be an unsustainably high number of people on the planet for as long as the fertility rate of Immortals was more than zero.

It might be the case, then, that Immortals could reasonably be expected to be sterilized. This would be a curtailment of reproductive freedom; but while many bioethicists hold to the idea that respect for autonomy is not just a principle that we should follow, but is inviolable, it is not obvious that they are correct. And, anyway: if a person is told that sterilization is a condition of Immortality and chooses to go ahead with Immortality anyway, then that is hardly a curb on her autonomy – any more than is being told by the man in the showroom that paying for the Lamborghini is a condition of driving off in it.

From the point of view of achieving the good life and maximal flourishing, it might be the case that not being able to have children would be undesirable; on some accounts, we might want to hold that reproduction is an aspect of the good life. However, this point is fairly easily dealt with. First, though having children may be a part of the conception of the good life entertained by *some* people, it is not part of the conception entertained by *all*: plenty of people believe that they are flourishing, or can flourish, without children. Moreover, there're ways to have children or a family that don't involve close biological

relationships: what benefits there are from family life do not seem to be reducible to any consideration about the lines of genetic descent. But even allowing that an argument was generated that established the contribution to flourishing made by having children of which one is genetically a parent – well, even then, it might simply be that there's a decision to be made about how to achieve the good life; genetic parenthood and Immortality may simply be antagonistic. Some may choose the former, and others the latter.

So it would seem to be, at the very least, less than obvious that the overpopulation argument will undermine the idea that Immortality could feature as a part of the good life.

## Exclusive immortality

But supposing, for the sake of the argument, that overpopulation would be a problem: one retort to this is that opting for Immortality does not imply that there would be any other Immortals. Wanting immortality for yourself doesn't mean that you want it for everyone else, too. Hence, if we are disposed to think that it is better to live for as long a life as possible, we need not be worried by the undesired consequences of everyone else living just as long. My flourishing might not be maximized in a world of generalized Immortality – but that tells us nothing about whether I would find it desirable for myself.

One interjection that might be made at this point would run along the lines that the quality of the life of a solo Immortal would be diminished simply by virtue of his solitude. It might be pointed out, for example, that a flourishing life is a social life, and that – perhaps more importantly – humans as a species tend to form deep and long-lasting emotional bonds with a few people. For these bonds to be broken – for a spouse to die, or a friendship to founder acrimoniously – can often be traumatic. An Immortal would therefore have to choose between sacrificing meaningful relations with others, or enduing a succession of bereavements. (Mortals have a reasonable chance of not having to experience the death of their loved ones because they also have a reasonable chance of dying first: a good way not to be widowed is to be the one who leaves a widow behind.) The alternative would be that all those about whom he cares should be Immortal, too – but those other Immortals would presumably have a circle of attachment that was non-congruent with the first's – and so we would very rapidly see ourselves having to accept a very large number of Immortals.

However, the nod towards friendship indicates the reason why this challenge is weak; friendships do come and go over the course of a lifetime – as, indeed, do other close relationships. Though these events do leave a mark on people for a while, they are rarely so devastating as to inhibit flourishing for the rest of one's life. Indeed, it's possible that we'd think that someone whose quality of life is devastated utterly by losing a loved one may be in some sense pathological – and there's certainly no reason to think that a person *shouldn't* live a full life in

the wake of such events. At the same time, though there is a reasonable chance that one will be bereaved by the death of a loved one at some point during one's life, this is not much of a reason to eschew love. Much the same could be said in respect of Immortals; they may find that their lives are marked by more of these events than are the lives of their mortal counterparts; but there is little reason to suppose that this will tell us that Immortality is a silly thing to pursue after all. An Immortal will know that there is a reasonable chance that she will lose those she loves; but this is not a reason not to love to begin with.

Nor is it quite clear why Williams thinks that an Immortal's personal relationships would require 'peculiar kinds of concealment' and would be vulnerable to 'a form of isolation'.[20] It is not obvious that an Immortal would be able to – or would even want to – avoid personal relationships with a mortal; after all, in the familiar world, people do form close relationships with others whose life expectancy is, for one reason or another, much shorter. Correspondingly, there is no clear reason why an Immortal would have to hide her immortality in order to have any chance of a meaningful relationship with a mortal. All relationships – whether with other people, or even with pets – bring with them the risk of their ending; but this is not a black mark against relationships.

There is no obvious reason to suppose that Immortality would be uniquely antagonistic to any of the social aspects of a good life. But none of this really answers the question of whether significant life extension would reasonably be desirable to someone interested in the good life to begin with. I am not convinced that it would be.

## The tedium of immortality

The beginnings of an argument against the desirability of Immortality are nevertheless suggested by Bernard Williams in 'The Makropoulos Case': that being an Immortal might well turn out to be intolerable. This being the case, there would be less of a reason to pursue Immortality, and it might well turn out that to insist on a *right* to it (and de Grey thinks that there is such a right[21]) would be irrational. In fact, mortals might have a significant advantage in respect of achieving the good life; and – though this is a touch more speculative – at least some Immortals might have a life that's noticeably worse, for reasons that I'll mention later.

Williams is happy enough to accept that dying is undesirable; but he is famously unhappy with the idea that not dying might be preferable. We need no reminding, he claims, that death can come too early; but the story of Elena Makropoulos reminds us that it can also come too late. Its tardiness is due to the fact that an unending life would become oppressively boring; and, as such, we would seem to be entitled to suppose that Immortality would be antagonistic to the good life.

Williams recognizes that life is good to the extent that it is required for the fulfilment of our desires, and I've already noted his claim that wanting something provides a reason for resisting things that exclude having that thing: things like dying.[22] The desirability of the life of the Immortal would depend on having a succession of things to do, careers to pursue, and so on, for the sake of keeping him occupied. But the problem that Williams identifies here is that the train of diversions would itself become tedious. After some particular time, one would have done and experienced all the things that it would make sense for a person of stable character to do and experience.[23]

It's possible that, if Williams is correct in this, things are even worse than he imagines. One reason for this is that, while many of us feel that it would be better to have more time in which to cram the things we want to do before we die, the Immortal would soon enough find herself looking for things to do with which the time could be filled. Life, that is, would become something that had to be endured. And yet the prospect that immortality would be something that has to be *endured* is, presumably, not what an Immortal would have wanted. And things get worse still, since not only would an Immortal have to find ever more ways to fill the time ahead of her, but this task would itself become an increasingly difficult chore.

One possible way of circumventing the problem would be to opt for a kind of 'elective amnesia', by which memories would fade over the course of, say, a century or two: this would mean that an Immortal would be able to rediscover some skill or pastime abandoned long ago as though for the first time.[24] There would on this account be a rolling disconnection between an Immortal and his distant past; but since what concerns an Immortal is his future (and the possible lack thereof) rather than his past, this ought not to be the sort of thing that worries him too much. This is a possibility that Williams does not seem to countenance: he spends time considering the question of whether a series of psychologically disjointed lives *heading into the future* could be an object of hope to an Immortal as a means of avoiding tedium, and concludes that there are difficulties with it[25]; but he doesn't seem to pay much attention to the possibility that a series of psychologically disjointed lives *heading into the past* would mean that experiences might never be felt to be repetitious. An Immortal wouldn't even have to opt for his memories being positively erased: all he needs is to be willing to countenance the possibility that they will become less and less complete as the centuries roll on – something that seems intuitively plausible. Hence there need be no shortage of novelty and fascination in an Immortal's life.

The potential fly in the ointment here is that a reasonably clever Immortal would quickly reach the conclusion that, since he cannot remember what happened more than (say) a couple of centuries ago, he will not be able to remember today a couple of centuries hence; and we might worry that this would rob him of any incentive to commit to any projects in the meantime

because their absurdity will be manifest to him. This is related to, but different from, the dread that might arise out of absurdity – the knowledge that, in all likelihood, nothing you do will matter in a thousand years' time, and it certainly won't matter in a million. The difference is, though, that most of us can still cling to the expectation that our actions will matter to us for the duration of our lives. This is not a hope available to the Immortal, since his past actions losing their salience is a condition of his continuing life being tolerable.

Still, there is a way to avoid this problem, and it is to insist that it is possible to pursue projects for the sake of the virtues they offer now, without reference to the extended future. Let them be absurd! Their absurdity really doesn't matter: what counts is that we should live for the moment. This is an impressive response – but, as we'll see later in the chapter, it also spells the end for the desirability of Immortality.

The part of the problem that remains unaddressed is that the badness of boredom has less to do with what has gone before than with what is yet to come: one can put up with almost any amount of boredom so long as it is in the past. The Immortal, though, knowing that he has an indefinitely large number of years spanning ahead of him, is susceptible to a kind of dread stemming from the expectation that the attractions of whatever he does do, no matter how enriching it is, will almost certainly pall; there will, then, be a never-ending need to find something new; and the difficulty of finding new things would rise exponentially. Knowing that you've done it all before is bad; but so is knowing that eventually you will have done it all. In this respect, Williams' concerns about disjointed future selves are well founded; for if my chosen method of avoiding this dread is to sever the connection between me and the entity that I will become such that I am no longer troubled by an indefinitely long life ahead of me, the sense in which I would still be willing my indefinite survival becomes unclear. For as long as the Immortal knows he is immortal, and projects a recognizable version of himself into an indefinitely far future, the dread of a predictably barren future will be likely to be inescapable.

## Self-inflicted boredom?

As a slight aside, it's worth noting that Immortals might turn out to be temperamentally unsuited to an Immortal life anyway, precisely because they will positively – and paradoxically – *avoid* certain ways to fill it. This is an argument that's articulated by Nicholas Agar, who speculates that anyone who takes seriously the idea of becoming an Immortal would lead a life that is rather restricted in other ways. Agar suggests that we might find that would-be Immortals lose out on a good portion of a good life by virtue of being rather risk-averse.

The argument is essentially this: For someone who thinks that more life is worth having in its own terms, anything that increases the risk of dying, either naturally or 'artificially' – say, in a road accident – would be seen as something to be avoided. This much is true for all of us, of course; but the difference is that if someone is willing to take extraordinary measures not to die from natural causes, it's hard to imagine anything but a correlative willingness to take extraordinary measures in respect of other causes of death.

> It's possible that negligibly senescent people will decide to treat the interior causes of death, such as the seven deadly things[26], differently from exterior causes of death such as driving, bush-walking, and skiing. The one-thousand-year estimate of the life expectancies provided by SENS is, after all, based exclusively on the elimination of the seven deadly things. It does not assume that we've done away with any of the mortal threats posed by cars and tsunamis. Is there anything wrong with pronouncing yourself content with the mere one-thousand-year expected increase?
>
> Though there's nothing wrong with saying this, I doubt that people susceptible to de Grey's rhetoric will be inclined to agree. It sounds suspiciously like the "pro-aging trance" that encourages us to accept the risks of death that come from the seven deadly things. Those who take de Grey's therapies have quite deliberately set themselves on a course of reducing the probability of death in succeeding units of time. . . . Negligible senescent people have shown themselves unsusceptible to the pro-aging trance in respect of the interior causes of death. They'll be spending much of their time self-administering anti-aging therapies. It's hard to imagine them cavalierly dismissing similar risks accompanying downhill skiing and helicopter flying.[27]

Immortal life, Agar worries, will be somewhat etiolated. It is not a given that such a life would be all that desirable; and so the quest to live longer might precipitate a life whose liver feels there is less to live for. An Immortal, inasmuch as that he's the sort of person who takes great care to avoid dying, may turn out to have to take great care not to do anything much at all – and that's hardly living a good life. (One thinks here of George Burns' joke that someone who gives up smoking and drinking doesn't live longer – it just feels like it.)

There is a part of this analysis that misses something. In a possible future in which the damage from injury or Kurzweil and Grossman's (rather ill-defined) toxins could be repaired or avoided easily, it might be possible for people to live rather *more* recklessly, in at least some respects, given that repairing the damage will become easier. But even taking this on board, Agar's claim does have the air of something correct about it as well. Admittedly, it's possible that Immortals will just learn to narrow their horizons. But, in that case – and this is a point that I don't actually see Agar making, though it's implicit – it's hard to see why people who're willing to narrow their horizons to accommodate the short life to which we're currently used are in any more of a trance than anyone else.

Either way, though, we're still confronted by people who appear to want life for life's sake. Either way, they're still going to have to find things with which to fill the extra time; the only question is how wide or narrow the range of acceptable options is. Either way, there'll come a point when all acceptable options are exhausted.

## Filling a life, and the LOT revisited

We can, though, put all these considerations about possible methods for avoiding boredom and angst to one side, because they miss the point. Attempting to find ways to fill a prolonged life is indicative of having misunderstood what it is that makes life worthwhile anyway. This becomes apparent when we revisit one more time Williams' claim that the badness of death is attributable to life being a necessary criterion for the realization of good things. What this means is that life *qua* life is valuable only in respect of other things. Put another way, death's badness is to do with it inevitably preventing the realization of projects. It's the projects, though – that is, the content of a life – that matter, rather than the fact that it's a life at all.

For the mortal, life is a necessary substrate for the pursuit and potential realization of a project (or the attainment of any other good thing that a life might contain); dying is bad because it hinders the things that really matter. This is basically the idea behind the LOT. But what is important is that the opposite is true for the Immortal: projects (and other good things) are not things that require life, but things that are required for his life to be bearable. They are simple props to occupy one's time. Instead of being the reason why one might feel that one wants more life, life becomes a reason to want more projects. That is at the centre of the Immortal dread of the future.

This might give us pause to wonder whether being an Immortal really is something to which it is sensible to aspire. However, if we reach the conclusion that being an Immortal would be bad after all, this still only leaves us in a position that was marked out by Williams: via the application of the LOT, dying is still bad; it's just that not dying is bad for different reasons. However, we can move things on a little from here, and ask whether we have to accept the truth of the claim that dying is really all that bad for the mortal after all. I am not sure that we do: I think that there is a good chance that a good life may be compatible with mortality after all. And if I'm right, it means that there's a way out of the problem with which Williams leaves us.

To get the argument going, it's worth recapping a little. I'm going to keep hold of the idea that flourishing for a human depends in some way on having and pursuing projects of one sort or another. One of the reasons for this is that a person's projects provide a shape to his life, and so provide us with a standard by which we – and he – can assess whether a life is a success or failure, which

might well shade our assessment of whether it is a good one. Another reason is that a person without projects would live a listless sort of life because he would have little or no stake in his own future; moreover, that even minimal projects are a necessary way of fending off boredom, and boredom is incompatible with a good life, means that someone interested in living the good life – *ex hypothesi*, this means everyone – will have an interest in pursuing one or more projects of one sort or another at almost all points in their life.[28]

I made another claim in Chapter 1, though, and it is very important for the position I want to outline here. That was that the good that is contributed to the good life by a project is, at the very least, not closely dependent on its completion. What is required to stave off boredom and to give a life a shape is that there are projects, not that there were projects that have since been realized.

There is a negative explanation for this, framed in terms of what the completion of a project does not do, and a positive one, framed in terms of what it does. The negative explanation is that while one might justly derive satisfaction from the completion of a project, that is not enough: and no matter how much satisfaction is generated, it will not itself be able to tide us through even an un-extended life. The positive explanation is that though it's the prospect of realizing a project that gives a shape to a life, it is the actual engagement with a project *right now* that eliminates boredom. Hence the contribution made to the goodness of a walker's life by hillwalking is attributable to a range of things: the anticipation of the walk, the actual walking of it, and so on. Having completed the walk is, of course, a source of satisfaction, and the project itself contains what we might call 'the idea of its completion': but its value does not come from its completion, and the goodness of having once completed a walk tarnishes reasonably quickly.

(By the claim that a project contains the idea of its completion, I mean to say that for a project not to be absurd, we must be able to imagine a world in which it is capable either of being completed, or in which progress towards its completion is possible. I find it hard to see how someone could truly be engaged in a project that she believes to be futile, and so think it to be worth the effort or the lifespan.[29] Note, though, that a person could be engaged in all manner of grand projects without absurdity. For example, it might be that a solution of the *P=NP?* problem is, as a matter of fact, beyond human mathematicians; but even so, for as long as we do not know this, it would not be absurd for a person to adopt solving the problem as a project (and so as a reason to be alive). A person could derive satisfaction from the project of painting the Forth Bridge, too – a task that is famously unending – because it is possible to imagine the whole bridge painted[30]; the job is not futile as a matter of principle, and so a person could be engaged in it. WG Sebald has an anecdote about a man engaged in the task of building a scale model of the Temple of Jerusalem in his barn – a demanding enough task in its own right, but one rendered

apparently interminable by the accretion of more archaeological discoveries and controversies, which require constant revision of the model[31]; but the Temple did look like *something*, and so we could imagine an accurate model emerging at some point.)

What is important here is that there can be projects that give our lives shape, but which we do not value for the sake of their completion so much as for the process by which they approach completion, and the completion of which is not required for their being valued. The fact that a project contains the idea of its completion generates the illusion that it's completion that matters; but, at least from the point of view of living a good life, it isn't. And, of course, if we *do* realize all the projects that we have at time *t*, it won't be long until we start to get bored again, and feel the need to find another project – so it might actually turn out that there could be situations in which a project's realization turns out to have been an undesirable thing. Immortals might find that they are condemned to finish what they start.

## Mortality and the good life

The nub of the hypothesis is indicated by a citation of Ovid in Montaigne's essay 'To Philosophise is to Lean How to Die'. Montaigne writes that '[w]e are born for action: *Cum moriar, medium solvare inter opus*.' The Latin he cites here is cribbed from book II of Ovid's *Amores*[32], and Montaigne takes it here fairly literally as 'When I die, may I be in the midst of my work'. For Montaigne,

> I want us to be doing things, prolonging life's duties as much as we can; I want Death to find me planting my cabbages, neither worrying about it nor the unfinished gardening.[33]

For the record, it has to be said that Montaigne misrepresents (and misquotes) the passage a touch (though in a way that leads more naturally to the point about planting cabbages): it's fairly clear from the context that Ovid is making a joke about the best kind of death being the death that one dies while having sex[34]; and it's stretching things a little to suppose that Montaigne means 'planting my cabbages' to be read euphemistically. But the general point stands: Montaigne's view is that, in being born for action, it is better for us that our lives end while we are actively involved in something that gives our lives both pleasure and meaning. A man who is content planting cabbages is, at that moment, living a good life; and for him to be so absorbed in that activity that (if I can be excused the image) he barely notices his own death has ended his life well. It doesn't matter if there's another row of cabbages left untended. Montaigne barely seems bothered by death at all, and certainly less concerned about putting off death than about the characteristics of the life being lived at the point of death.[35]

This sentiment is sent up wonderfully in Joshua Ferris' novel *Then We Came to the End*:

> Hank Neary had a quote and we told him politely to shove the quote up his ass. "When death comes, let it find me at my work." He said he couldn't remember if it was Ovid or Horace who said that and we replied that we could give a good goddam what Ovid the Horse said. Ovid the Horse got it wrong about death and work. We wanted to die on a boat. We wanted to die on an island, or in a log cabin on a mountainside, or on a ten-acre farm with an open window and a gentle breeze.[36]

But it turns out that Ovid the Horse, Montaigne and the narrators of the novel are more or less of one mind. For all, what counts is that when death finds them, it should find them content with whatever it is they are doing. For some, contentment is found in planting cabbages; for others, it is found in a log cabin. I have argued that a good life depends, for the most part, on doing something. I find it unlikely that a sustainedly good life can be had simply by virtue of being on a boat, since a good life demands engagement and simply being on a boat does not (as far as I can see) provide that. But that might just be because I've missed something about boating. All the same, what matters is engagement: with living life in a manner that promotes flourishing – whatever that manner may turn out to be. If a person finds satisfaction in work of a certain sort, then doing that work may be a way to live a good life. This is a fairly egoistic position – but the quest for the good life is, in a sense, irreducibly egoistic to begin with, and no worse for that. The point is that, if one's life at a given point is enriched by what one is doing, then it is – at least at that moment – being lived well; and a life that ends 'on a high', while it's being lived well, ends better than one that doesn't. It makes sense to hope that death finds us at our labour.

Williams seems to be nonplussed by the prospect of this kind of absorption in a project. In response to the notion that Immortality would be bearable if we could find something endlessly and totally absorbing, he claims that

> if one is totally and perpetually absorbed in such activity, and loses oneself in it, then as those words suggest, we come back to the problem of satisfying the conditions that it should be me who lives forever, and that the eternal life should be in prospect of some interest.[37]

To some extent, Williams appears correct. If I am willing to lose myself in a task, I am willing to lose myself – in which case, I lose contact with who- or whatever it is that goes on to live forever. Correspondingly, if I am concerned that *I* should live forever, this seems to be inseparable from a picture of my continued self; I ought to be able, to some extent, to adopt the pose of a spectator of my own activity. But what Williams does not account for is that, to the extent to that one is a spectator of one's activity, one is alienated from oneself. To be absorbed in one's projects is to be devoted to something

that – presumably – one thinks worthwhile; to stand as a spectator to that kind of absorption is to deprive oneself of whatever satisfaction it offers. Put another way, a good life is one in which we can lose ourselves in the things that we think desirable – whether they be sex or cabbage cultivation.

If I'm right to suggest that it's engagement with an ongoing project that matters, and that completing it doesn't – or, at most, matters only a very small amount, since the satisfaction that one derives from having completed a project is transient, having a tendency to drift over the horizon – then we have another reason to suppose that extending our lives isn't all that important after all. The fact that we need to be alive to complete a project means that we need to be alive only if the completion of that project is important to our lives being tolerable – but that is not true. The goodness of a good life is related to the completion of projects in, at most, a glancing way. Nor is it obviously better to have time to complete a number of projects. If something, or some range of activities, is engaging now, then that is all that is needed for the life that we currently live to be rewarding. We may have a simmering hope that we will have engaging activities in the future – but this is a version of the hope that, if we are alive, our life will be rewarding, and it won't therefore tell us anything about the desirability of more life *per se*. Naturally, we would subscribe to the idea that, *if* we are alive, our lives should be good; but this is not the same as subscribing to the idea that we ought to remain alive *so that* that life can be good. All we need to live a good life is that that life should be engaging for its duration; duration doesn't add to the goodness.

## The boon of mortality

What I hope to have done in this chapter is to deflate the case for immortality as being a serious contributor to the good life. Immortality brings with it the risk of boredom, and – more importantly – it makes the mistake of thinking that more life is desirable just by virtue of there being more of it, as though life is rewarding quite aside from anything that goes on during it. Once we get hold if the idea that it's the content of a life that matters, though, it's only a short step to being able to suggest that living longer adds nothing except time that needs to be filled.

Accepting this means that we can confront de Grey's claims about being caught in a pro-ageing trance, too. For it is possible for people to live a good life notwithstanding their being liable to ageing; and accepting and adapting to ageing is hardly the same as adopting a *pro-* attitude. It might yet turn out that all these apparently happy people are deeply mistaken, and not happy at all; I have no principled objection to this kind of bad faith being possible. But, absent more evidence, it is equally possible simply to accept that elderly people who claim to be basically happy really are basically happy. This happiness is

frequently to be found even in those who have chronic health conditions – people who have been disabled by stroke, or who have diabetes or whatever, but who manage their lives within constraints. (Maybe we'd want to say that they shouldn't have to adjust to the facts of ageing; but not *having* to is a long way from it being an indication of bad faith.) A lot of this has to do with a wide, varied and supportive social network. The goodness of these lives comes, that is, not from having got rid of senescence – but from having got rid of the *worry* about senescence. That, too, is something that seems to have been forgotten by the Immortalists – their picture of human life is not only a picture of an empty box into which things can and must be crammed, but also one that is individualistic and isolated. If one is socially integrated, though, growing older is just one characteristic of life; a part of facticity that has to be confronted, and nothing more. It's even possible that there's something slightly petulant about the Immortalist's view that his life is worth the effort of saving. Really, it isn't; and I have a hunch that there's something quite admirable about the cynic who recognizes that he should chisel what he can from his life, knowing that it counts for nothing much over all.

There's one final aspect to the defence of mortality, and it harks back both to Agar's concerns about the risk-averse Immortal, and to Montaigne's comment that he wants death to find him planting his cabbages, 'neither worrying about it nor the unfinished gardening'. The final thought is this: Immortalists worry too much. In thinking about ways not to die, they are inevitably thinking about their own death, and so about something they think to be bad. Their deaths cast shadows over their lives. Montaigne thinks about his garden, about the butterfly on the flower, about his wife and mistress, and about the sunshine. I'd rather be Montaigne.

## Notes

**1** Harris, 2007, p. 61.
**2** Ibid.
**3** de Grey, 2005, p. 662.
**4** Harris, 2007, p. 52.
**5** de Grey, 2005, p. 659.
**6** Seneca, 2010, 4.3.
**7** Harris, 2007, p. 64.
**8** Finding a vaccine for HIV has proved to be difficult; but there are several strategies being pursued – see, for example, work by Balasz et al. (2012), among many, many others.
**9** Kurzweil and Grossman, 2011, p. xxiv.
**10** Ibid., pp. xxiv–xxv; slightly modified.
**11** Harris, 2007, p. 68.

**12** I'll use the capitalized form 'Immortal' here to designate a person who, thanks to some scientific intervention, can have a reasonable expectation of a life several times longer than the current average.
**13** Nagel, 1970, p. 80.
**14** Ibid.
**15** Williams, 1999, p. 85; slightly modified.
**16** Adler, 2010, pp. 38–9.
**17** Kurzweil and Grossman, 2011, pp. 207–312.
**18** Ibid., p. 208.
**19** This point is obviously borrowed from Parfit's *Reasons and Persons*.
**20** Williams, 1999, p. 90.
**21** de Grey, 2005, *passim*.
**22** Williams, 1999, p. 85.
**23** Ibid., p. 90.
**24** In Chapter 6, I consider the possibility of mnemonic enhancement – an intervention designed to increase a person's capacity for memory. It occurs to me in passing that such enhancement may seem like the opposite to a bored Immortal.
**25** Ibid., *passim*.
**26** The 'seven deadly things' are the seven biological processes that de Grey holds to be the causes of senescence.
**27** Agar, 2010, p. 120.
**28** Almost all, but perhaps not all. A person might quite straightforwardly have space in his life for a degree of empty-headed sitting on a beach for a few days each year; this is compatible with avoiding self-stupefaction. Still: it is reasonable to think that this would have its limits.
**29** Albert Camus suggests that we ought to imagine Sisyphus happy – but this is a struggle almost as great as Sisyphus' own; it seems likely that Sisyphus would have to put rather a lot of effort into believing in his happiness.
**30** As I was drafting this chapter, the BBC reported that the painting of the bridge had, in fact, come to an end for the time being, which just goes to show, doesn't it?
**31** Sebald, 2002, p. 242.
**32** Ovid, 1982, p. 124.
**33** In Montaigne, 1991, p. 99.
**34** Indeed, Peter Green's translation of Ovid's poem translates 'inter opus' as 'on the job', which seems much better.
**35** For more on Montaigne's thoughts about his own death, see Chapter 1 of Sarah Bakewell's *How to Live*.
**36** Ferris, 2008, p. 327.
**37** Williams, *op cit*, p. 96.

# 5

# Designs for Life

I argued in the last chapter that it's not clear that the indefinite postponement of death, though probably perfectly morally permissible, would have to make a great difference to the goodness of our lives; it is therefore possible to doubt whether it's something that would be high on the agenda of someone who was interested in maximizing his flourishing.

The intention of this chapter is to defend the idea that there are some enhancements that are potentially worth defending and pursuing. I shall claim that there is a range of modifications and manipulations of the human body that can much more easily be justified. And there is a couple of notable, and related, characteristics of the justification. The first is that the welfare component plays only a small role – and in fact, it might be that some of the modifications and enhancements I have in mind could be construed as being contrary to welfare. The second is that at least some of the modifications I have in mind are more-or-less useless, and so might be expected to rank as being among the least worthwhile things that we can or should do from some perspectives.

I'll begin by considering the use of enhancement technologies in sport. This part of the chapter will not address the contribution that enhancements can make to the good life directly, but it is useful for a couple of reasons. First, sport provides a context in which there is already a great deal of discussion of enhancements (and a great deal of this discussion is not confined to academic debate) – but it is a context in which the positive arguments in favour of enhancement seem to be tolerably, and perhaps surprisingly, strong. Second, the discussion of enhancement in sport will lay the groundwork for what comes in the subsequent parts of this chapter. The enhancements that are first to cross most people's mind in respect of sport are pharmacological – but in the middle third of the chapter, I'll push the boundary of this a little, and look at physiological enhancements. This will provide a bridge to looking at other interventions that a person might make to the fabric of the body – and these do, via a version of the claim about projects that I've been pushing so far, potentially contribute to the good life.

## Enhancement in sport

There is a number of contexts in which a person might feel that they have a reason to make use of enhancements. Sometimes these reasons will be to do with gaining positional advantage; at other times, they'll be related to non-positional benefits. One of the most obvious contexts is provided in the sporting arena, where athletes may seek to improve their performance by chemical or biomedical means. Indeed, 'enhancement' is a word that is very easily associated with sport, and any number of scandals and disqualifications of athletes found to have used performance-enhancing techniques of one sort or another.

There is a number of ways that an athlete (or his performance) could be enhanced, and the most obvious is by doping: the administration of some pharmacological agent that boosts muscle development. However, athletes might also consider 'blood-doping', which is the practice of siphoning off red blood cells and then re-injecting them to boost oxygenation capacity; and it is possible that we might one day see gene-doping, which would involve parents selecting or implanting genes in their unborn that will make them more likely to be high achievers. Whatever the technique, though, a great deal of effort is devoted to ensuring that athletes are not doped; and there is no shortage of public interest or scandal if an athlete is found to have used performance enhancers.

Is there any justification for this, though?

For the time being, I'll limit my attention to pharmacological interventions – although other kinds of intervention would probably not be too hard to work into the argument, and what goes for pharmacological interventions probably goes for them, too. There's a range of arguments that we might find marshalled against the use of enhancements in sport. One is that to use enhancers is cheating, and – since we ought not to cheat – we ought therefore to avoid them. A second is that there are safety worries that merit a ban on enhancements. A third is that the use of enhancement erodes desirable characteristics in athletes; and a final concern is that performance enhancers erode the character of the sport in question. As we'll see, it's only the last one of these that is particularly ethically interesting.

The cheating argument can be eliminated fairly quickly and easily. Though I have no desire to deny that cheating is wrong and that athletes ought not to cheat, the premise that performance enhancers are necessarily cheating is patently false. Quite obviously, athletes make use of all manner of performance enhancers. Participants in the newly-revived Olympic games at the end of the nineteenth and beginning of the twentieth century took strychnine, for example; modern athletes take full advantage of the fruits of dietary science to make sure that they get the optimum balance of nutrients in their diet. Moreover, they make use of shoes designed specifically for their discipline; swimmers wear

goggles so that they can see what they're doing, and cyclists wear helmets that allow them to take greater risks when overtaking. Athletes train at altitude; and they avail themselves of all manner of other techniques and equipment to make sure that their performance is as good as it possibly can be on the day of the competition. Since all these things are undeniably intended to be performance enhancing – that is the whole point of them – and at least some of them will be effective, and since none of these things normally raises an eyebrow, we have to admit that performance enhancement *per se* is not a problem; and nor is it cheating, because all these things are allowed. What *is* a problem is illicit enhancement; and this is a problem because it tilts the playing field in favour of those willing to break what rules there are, and is therefore unfair. However, this will not tell us that the rules are necessarily all that morally compelling, or that anything would be lost were they to be altered. Were the rules altered, using (say) steroids would no longer be unfair, since they would simply be one more technique that an athlete could take or leave. More generally, it looks so far as though the use of performance enhancers, if it is wrong, is *malum prohibitum* rather than *malum in se*.

But – we might worry – this leads naturally to the safety objection. Mightn't the worry that his competitors are using it drive an athlete to use a drug the safety of which is not clear? Many drugs are unsafe in at least some sense; the justification for using them therapeutically is that the dangers they present are outweighed by the medical benefits. But for the athlete, there are no medical benefits, because there is no illness; and so the warrant for taking the risk may seem to be accordingly less. However, this argument has to overcome a number of objections. First, it has to demonstrate that medical benefit is the only thing that can warrant taking pharmacological risk, and this is far from a given. In Chapter 1 I suggested that ill-health (and, implicitly, health) could be understood within the context of a set of claims about projects and the possibility of their fulfilment; and if this is correct, it seems to indicate that it's those projects that are in the moral driving seat and could help determine whether a certain procedure or intervention is warranted: medical benefit is, on this account, at least partially a function of other things. It might be that athletic glory is just one of the projects that a person might adopt. Whether a doped athlete can claim glory is a different matter, and one to which I'll return in a little while. But what counts at the moment is that, while illness may be sufficient to generate a warrant for intervention, we can't assume that it's necessary.

But that aside, there is no particular reason to suppose that biomedical enhancements *have* to be unsafe – and there's plenty of room to turn the safety argument against itself by pointing out that the prohibition of certain drugs may exacerbate safety worries by making quality-control and regulation impossible. Most importantly, if people were primarily concerned about their health, this would militate against elite athletics. Though sport may have health

benefits for most of us, those health benefits are fairly low-hanging fruit; the kind of commitment required to become an elite athlete is far greater than that required for health. Given the injury risks and increased likelihood of (say) joint problems in later life, someone who was concerned to maintain their physical welfare probably ought positively to avoid elite athletics. This is as much as to say that those who want to be elite athletes are probably committed to taking risks with their health anyway – and the risk from a performance-enhancing drug is just one from a list.

It's perhaps worth adding one other thing here about the way the safety argument can be turned against itself: even if some drug has a history of being physically damaging to athletes, this might actually generate a reason to legitimize it. After all, though many drugs will have been tested on healthy volunteers for the sake of assessing their safety, not all will – there may be some that we expect to be risky and whose use is justified only by the greater risk of the illness they are meant to cure (one would not test radiotherapy on healthy volunteers, for example); and even among those that are tested on healthy volunteers, it may be that the physiological differences between the man on the street and the elite athlete would make a difference to the way that the body responds to the drug. Hence there may be a gap in our knowledge about safety in respect of a particular drug. The way to plug this would be by means of more research into how that drug is processed by the kind of body that we might expect an elite athlete to have – but that's exactly the kind of research in which an athlete has a strong incentive not to participate, precisely because it'd preclude competition. This is as much as to say as that keeping the drug banned might not be justified out of concerns for safety – it's being banned might be one way to ensure that safety worries cannot be properly addressed.

A wholly separate line of argument against the use of performance enhancers is that they erode some important characteristic that is worth preserving in the athlete. This is an argument advanced by Michael Austin.[1] For him, one of the good things about sport is that it instils in us admirable characteristics like self-discipline – for example, sporting excellence requires that we be prepared to sacrifice that extra beer or bacon sandwich for the sake of the 'higher' goal of sporting excellence. So even if we are not inclined to think that sport is a good in its own terms, we have a reason to accept that it is good inasmuch as that it nurtures attitudes that are worth nurturing for the sake of excellence more widely construed.[2] But performance enhancers, the argument goes, reduce the need for such desirable traits as self-discipline, and so run the risk of being morally corrosive. (Austin adds to this a claim that admirable people seek excellence in all fields, but that such moral corrosion would be the kind of thing that those who are truly excellent would avoid.) We might add to this the thought that the use of a performance enhancer is simply *infra dig*, because it allows an athlete to get something

for nothing: a cyclist cannot claim to be excellent if he's taken a drug for the same reason that he can't claim to be excellent if he's fixed a motor to his wheels.[3] This kind of argument might even speak directly to concerns about the good life. Martha Nussbaum's 'capabilities approach' to understanding human flourishing and dignity allows for us to make claims about 'central human capabilities' that are worth pursuing or nurturing.[4] The capability to be vicious is not one of the things worth cultivating – but, by the same token, we might think that the capability to display certain virtues *is*; and so we might try to encourage self-discipline not as a means to achieving the good life, but as a part of it.

Yet this objection to the use of enhancements fails simply because it overestimates wildly what the use of most performance enhancing drugs can do. An athlete may find that the use of certain drugs sharpens his performance, his speed of recovery from injury, his effectiveness in training, and so on – but unless he is prepared to make efforts that to the rest of us seem outlandish to hone his own performance, and unless he has the physical capacity to begin with, they will not make much difference. Notwithstanding that using a certain drug may allow an elite athlete to shave a tenth of a second here and there from his time, we ought not to lose sight of the fact that he would remain an elite athlete without it. A doped elite athlete is much superior to most of us in his chosen field – but so is an un-doped elite athlete, and the difference is marginal. Granted this, such techniques take none of the graft out of becoming an elite sportsman, and cannot compensate for the fact that most of us simply don't stand a chance of getting anywhere near an Olympic podium; they are aids to, not substitutes for, very hard training. So whatever the morally desirable aspects of the pursuit of sporting excellence, it doesn't seem at all obvious that it's threatened by the use of performance enhancers.

Relatedly, neither does there seem to be a great deal of mileage in the idea that the achievements of a doped athlete are not 'real', or *as* real, as the achievements of the un-doped, with the implication that doping removes from sport whatever it is that we find valuable. As we've seen, there's any number of performance-enhancing techniques already in use – things like controlled diet, hyperbaric chambers, physiotherapy, gait analysis and so on – and none of them strikes us as devaluing real achievement. It seems just as likely that we could say that these are things designed to cultivate real achievement. But it's not clear why we might not say the same about currently illicit performance enhancers: there's no reason in principle why we couldn't say that these aren't just things that help a performance be all that it can be. At least on the face of it, to insist that doping means that an athlete won't show real achievement is a bit like saying that someone who feeds his plants will not get real tomatoes. Call this the *Real Tomato Argument*; it'll show itself again in Chapter 7. For the time being, I'll move on to what is, I think, a better way to think about what is and is not acceptable in sporting enhancement.

## The character of the sport

So far, it looks as though enhancement *qua* enhancement is not too big a moral problem in respect of sport, not least because many enhancements are in everyday use and allowed by the rules anyway. For as long as they are legally accessible by all participants within a competition, each of whom can take or leave them as part of their preferred regimen of preparation, there is nothing unfair about them. Legitimization would also ensure that, to the greatest degree possible, justice could be maintained in the sphere of competition, since noone would be rewarded for breaking the rules as they may unwittingly be now. And safety would, likely as not, get better.

And yet this is not an encomium for performance enhancers. There is more to be said about the use of enhancers is in relation to what their use might do to the character of the sport in question: I want here to advance the claim that there is such a thing as one enhancement too many, and that this might give competitors a reason not to use at least some. All being well, this part of the argument will lead into a more general consideration of body-modification and the contribution that it may make to a good life.

The basis of the claim here is that there is room for a kind of essentialism in sport; there are some things that athletes cannot do while remaining true to the essence or character of the sport in which they are ostensibly taking part. Over-enhancement might mean that one has departed from the essence of the sport in question; hence there are certain enhancements that one ought not to use because they are antagonistic to the end that they're meant to achieve. The idea here is perhaps best served by giving a couple of examples.

I mentioned almost as a throwaway a moment ago that a cyclist who fixed a motor to her bike would struggle to maintain the claim that she was a particularly great cyclist. But it's actually quite a significant point: in having attached the motor, she has departed from what a cyclist is generally understood as being. As such, she would not have enhanced her capabilities *qua* cyclist. Similarly, there might be room to wonder whether an athlete who has had his legs amputated and replaced by carbon-fibre blades is an enhanced sprinter, or something else. (Oscar Pistorius, the South African blade-runner, competed in the 2012 London Olympics, but his participation was not without comment.) This sort of consideration does not apply to people who make use of hyperbaric chambers, or even pharmacological enhancements – whether or not we approve of the method used to improve performance, the nature of the activity in which they are engaged is unchanged by their attempts at enhancement.

Of course, where we draw the boundary between what is within the character of a sport and what turns it in to something else is as tricky to decide as where we draw the line between species as one evolves into another. Sometimes an enhancement might not detract from, but might even contribute to, the spirit

of the event: as Sandel admits, 'running shoes highlight rather than obscure the excellence the race is meant to display'.[5] At other times, it might be genuinely difficult to decide whether the line has been crossed. However, this does not mean that there is nothing to say. Those with an acknowledged expertize in the sport – either as competitors, pundits or fans – will be able to draw on that expertize to enter meaningfully into debates about what it is reasonable to expect from athletes and what it is reasonable to forbid. By analogy, though there might be no litmus test concerning what makes a good gymnastic display, there is still such a group of people who can claim expertize, and who constitute the judging panel; the same sort of process applies here. Though there may be no bright line dividing what is and what isn't acceptable, this doesn't mean that there can be no reasoned expert debate on the matter.

It's the appeal to expertize of this sort that helps us solve the normative question raised by the appeal to a given sport's essential characteristics; these will be defined by a set of features that are agreed by people with a stake in that sport to be vitally important to its nature. The assessment may well be irreducibly aesthetic, and perhaps not defensible by terms that can easily be articulated: it might just be that there are some performance enhancements that are held to be *just not cricket*. The standards agreed by the experts in question would provide a means to justify the current rules at a given time; they would not need justification in themselves.

This kind of analysis would seem to fit well with arguments placed against certain performance-enhancing techniques by real sportsmen. For example, there exist kinds of swimsuit made from polyurethane, which minimize drag within the water. One might expect these to be popular with swimmers; but this is not actually the case. The BBC reported that the Olympic swimmer Michael Phelps supported a change in the rules to ban them in 2009, saying that

> a solution needs to be found to the issue of swimsuits as *this is not swimming any more*. Technology has to go forward but I am looking forward to the time when we can go back to pure swimming.[6]

In one sense, Phelps' statement is nonsensical: swimmers using the new suits are, plainly, getting from one end of the pool under their own power – and, as such, it's hard to see how they are doing anything but swimming. But in a wider sense, Phelps seems to have been making a broadly essentialist claim: that the suits somehow eroded whatever it was that makes swimming the sport it is – hence the claim about 'pure' swimming. Those who use them may benefit from their effect; but if everyone can make a choice to use them, that benefit will cancel out. Phelps' point seems to be that it doesn't matter that it would, though: their use would erode the sport, and allowing everyone to use these suits would simply increase the speed of the erosion.

But even given all of this, if a particular performance-enhancing technique is deemed to be contrary to the character of this or that sport, it doesn't follow

that it is contrary to the things we find entertaining or admirable about sport in general. It might just be that we end up admitting that (say) running with carbon-fibre 'blades' in place of missing lower legs in the manner of certain athletes is a different sort of thing from conventional running, requiring a different set of skills, and something at which a person may be more or less excellent and admirable depending on speed, endurance, style, and all the rest of it.

## Becoming a blade-runner

But this does prompt a question. Suppose someone is watching an athletics meeting, and decides that they fancy having a go at competitive running, with a view to becoming the champion. The way is open for them to find an appropriate club, begin the training, and see what happens. But now suppose that someone sees Pistorius in action, and comes to believe that what he does is not running as conventionally understood, but is something different – that blade-running is *sui generis* as a sport; and so taken is he by this sport that he decides that he would like to be a blade-runner; and not only decides that, but adopts such a vision of himself as being central to his life, even though he has (at present) perfectly and unremarkably serviceable legs. How should we think about the possibility that he might want to take the step necessary to becoming a blade-runner, and have those perfectly and unremarkably serviceable legs removed?

Inasmuch as such an operation would be non-therapeutic, but would be carried out for the sake of realizing the would-be amputee's vision of a good life, I think that it would not be unreasonable to treat it at least *prima facie* as an enhancement – a surprising one, perhaps, but an enhancement nonetheless. Would it be permissible, then, or even reasonable, to amputate? As I argued in Chapter 1, the fact that someone *thinks* that something would be a good idea is not sufficient to show that it actually is; and a person may be mistaken in her identification of flourishing and the good life, even for herself. Might an amputation, then, turn out to be nothing more than a *prima facie* enhancement?

I want to argue here that there is no particular problem in principle with this sort of intervention. In doing so, I'll open the way to being able to make a range of claims about the way in which the biosciences broadly construed can make a significant contribution to our living flourishing lives.

In July 2010, a story appeared in the *Manchester Evening News* concerning a schoolgirl called Danielle Bradshaw. Bradshaw was, according to the article, born with a congenital dislocation of both hips and her right knee. Her preferred course of action was to have the right leg amputated. Amputation was not medically indicated – other surgical interventions would have allowed her to

keep the limb, although it would not have been of any use. However, from Bradshaw's point of view, another factor played into her decision. According to the news report, she decided to pursue amputation

> because she wants to become a world-class athlete. Danielle Bradshaw, 11, does not need to have her useless right leg removed but has chosen to go under the knife in a bid for sporting glory as a Paralympian.
>
> Tired of repeated trips to hospital and having to watch her friends play out without her, and inspired by gold medals, Danielle stunned her parents by declaring she wants to have it amputated. They agreed and she will have the operation in Sheffield in August. When the wound heals she will be fitted with a prosthetic replacement and friends are already rallying round to raise money for a running blade and sports wheelchair.
>
> She said: "I'm not scared, I'm excited. I can't wait for it to be done so I can start running and training."
>
> She said that seeing documentaries about soldiers who had lost limbs overseas and watching Heather Mills on *Dancing On Ice* helped her make her decision. She said: "I saw that and it helped make my mind up. And I am sick of all the operations. I just want to be a normal kid. I see people running and I want to know how that feels."[7]

It was reported in August 2010 that Bradshaw had had the amputation.[8]

There is a number of questions that we might be tempted to ask about this story. Some of them are fairly workaday in respect of the clinical ethics questions raised. For example, we might want to know who paid for the operation: noting that there is a difference between consent to and demand for treatment, there might well be a question to be answered about whether public funds ought to be used by the NHS to provide things that are medically unnecessary. We might wonder whether an 11-year-old really is capable of making a decision as big as whether to have a leg amputated; and, if she is not, whether anyone in fact would be. Finally, there is a question concerning whether the operation was warranted in any sense.

I'll put the funding question to one side, because it's peripheral to my concerns here. The question of warrant is much more important, and in order to get at that, and for the sake of the claim I'm going to be advancing, I'm going to tweak the *MEN*'s story a little bit. In my version, the protagonist of the story – let's call him Paddy – is someone who does not have any prior medical condition; hence the desire for amputation is wholly on 'lifestyle' grounds. Moreover, while someone who was really intent on becoming an Olympian blade-runner may very well be best off starting to train quite early in life, perhaps before the age at which we would be quite comfortable allowing them to make decisions about whether to have parts of their body chopped off, I'm going to allow myself the fiction that age does not matter – either all young teenagers are perfectly competent to decide they want a healthy and functional limb removed, or else older people will have no difficulty making up the ground in training.

To get the argument proper going, it's worth spending a little while looking at the justification for any surgery whatsoever; this done, we can apply this to the particular kind of surgery that would be required to remove someone's legs for the sake of becoming a blade-runner.

In most cases, providing the justification for surgery will be a very straightforward matter. Surgery is harmful and risky – it involves causing injury to the body and exposure to dangerous chemicals – but there is nothing too *outré* about the claim that inflicting this harm and this risk can be warranted, and frequently is warranted if it's a necessary part of the provision of a person's physical (or, in the case of at least some cosmetic surgery, psychological) well-being. This kind of rule offers us a way to say if and when surgical intervention is warranted – and nothing appears to be ruled out in principle. Thus even a major harm such as that caused by having your thoracic cavity opened will be warranted if it is a necessary precursor to the removal of a life-endangering tumour, for example. Neither does it have to be the case that the person who faces the risk from surgery will be the one who stands to benefit from it: extracting bone marrow is an invasive procedure, and brings with it a certain level of risk; but we generally do not have any trouble accepting it if the intention is to provide a treatment for leukaemia.

Working the other way, if we can say that a given intervention is warranted, this suggests that we could say in principle that another intervention is not. For example, there is a significant debate about the propriety of amputating apparently healthy limbs in the context of body dysmorphic conditions. Sometimes, we think that surgery should not happen.

When it comes to deciding whether the possibility of becoming a blade-runner is sufficient warrant for surgery on Paddy, it would seem that there is room for a great deal of disagreement. It is easy to say that there is a moral import to surgery that will provide medical benefits to the person undergoing it, or to another person; it is much less easy to say that those benefits are there for the taking in cases like this. It may be that no decent surgeon would be willing to operate. After all, if a would-be pole-vaulter decides that she doesn't enjoy the sport after all, she can return to her old life having lost nothing more than the price of some equipment and a few evenings at the stadium; this would be a waste of resources to some degree, but she would not subsequently find herself in a seriously harmed state. A would-be blade-runner, though, would have lost far more, and the amputation of the limb could potentially come to be seen as a significant and lasting harm to someone like Paddy. (Of course, the same thing applies in respect of other body modifications like tattoos – but these are less radical, can often be hidden, and can at least sometimes be erased completely. Something like scarification is not so easy to hide – but even this is much less radical than amputation.[9]) So maybe surgeons would hesitate before operating – especially if they think that their job is the promotion and safeguarding of health. And maybe they would hesitate for an indefinitely long time.

However – as should be clear from preceding arguments – there is no particular reason to think that health is the only thing that could, as a matter of principle, warrant surgical intervention. Being reasonably capable to engage with reasonable projects matters. Health *per se* is valuable inasmuch as that it impacts on those projects – but health without projects, howsoever minimal, does nothing to ensure flourishing; and it (or, at least, its importance) is best understood within the context of those projects and a life of a particular sort. Health, in short, is and ought only to be the slave of flourishing in a wider sense. If it's projects that matter, though – and, notwithstanding that a deep and abiding desire to be a blade-runner is perhaps a somewhat idiosyncratic project, though we might not agree that this ambition is all that pressing, it does not seem to be incoherent or any less reasonable than any number of other ambitions or projects that a person might adopt – then the point that the amputation of someone's legs does not further their health need be neither here nor there. If a surgeon thinks that what he does is enable people to realize their projects by means of therapeutic intervention, he might just as easily think that what he does is enable people to realize their projects by means of non-therapeutic intervention. Hence, though some surgeons may hesitate, others may be happier to operate without obviously having made an evaluative mistake.

There is, I would contend, no particular reason to suppose that the amputation of Paddy's healthy limbs and replacement with carbon-fibre blades would be particularly impermissible; and it is at least conceivable that Paddy would find himself flourishing more after the operation than before, for just so long as his ambition to be a blade-runner was undiminished. A non-therapeutic amputation could, in at least some cases, not only be compatible with flourishing: it might be something that's positively worth considering for the sake of maximizing it.

Of course, we ought not to dismiss easily the possibility that Paddy is about to make a big mistake that he'll come to regret about the kind of life that will generate flourishing for him or the means to achieve it; but there is at least a passing resemblance between his situation and that of any number of other people who might seek surgery for non-therapeutic reasons. A man who has recently converted to Islam or Judaism will have to be circumcised, for example: there are situations in which it is fairly straightforward for a person to undergo some kind of surgery in order to realize a particular vision of the good life, even if it's a decision that's non-reversible in the event of regret. It's part of being a self-controlling adult – and that's as much as to say that it's a part of human flourishing – that we should be able to make decisions even if we might come to regret them. Our blade-runner seems to be at least potentially nothing more than a particularly unusual member of this class of people: he has a vision of the kind of life that he wants to live, and is willing to adapt his body accordingly.

Another possible objection to Paddy's proposed course of action is that there is some morally questionable aspect to any project in which a person seeks to make himself disabled. This thought may be cashed out in one of a number of ways. The first would be that the biosciences ought to aim to remove disability, not inflict it. An initial response to this is that it runs afoul of the social model of disability, according to which disability arises when the world fails to accommodate the needs of individuals: that the objection takes for granted a particular and ultimately disputable normative account of what bodies should be like. The most we can say is that the role of the biosciences is to help us live a good life, which is exactly what Paddy has in mind. Moreover, if the point of the biosciences is that they should act as enablers rather than disablers, Paddy seems to be able to say that the replacement of his limbs with carbon-fibre blades *just is* an enabling intervention, inasmuch as that it will mean that he can subsequently do something to which he attaches importance that he could not have done before. Nor would a slightly more sophisticated version of the objection – in which the word 'hardship' replaces the word 'disability' – carry all that much more weight. This is because Paddy could conceivably deny that the operation would cause hardship, or, if it did, he could insist that it was worthwhile. Training to compete as an elite athlete in more familiar sports also involves hardship, after all, and a comparison might be drawn with conventionally therapeutic surgery, which often causes at least some degree of hardship that it is justified by an appeal to the things that, overall, it will allow. The same sort of analysis, he might argue, could be brought to bear on his situation, which is not paradigmatically different.

Another line of objection may be that there's something distasteful about an able-bodied person treating disability as a kind of consumer choice – that it diminishes the lived experience of the 'genuinely' disabled. However, this line of argument would, to be consistent, have to object to the choices of those who opt for surgery to remove disabilities when such surgery is possible; they, too, are treating disability as something that it's within their gift to accept or reject. And since that kind of objection looks to be absurd, it is not obvious that it would be any less absurd when applied to Paddy.

Finally, there may be an objection to Paddy's project along the lines that it is perfectly possible to live a good life without the intervention, and that there is a virtue to be found in accepting the life that one has, with all its limitations. Maybe Paddy ought, on this sort of line, simply to reconcile himself to the facticity of the body that he has. Such a line of argument might be derived from the equanimity in the face of death that I championed in the last chapter. However, part of the argument there was that life *qua* life wasn't worth pursuing except as a substrate of projects, which are. And although it oughtn't to matter to us whether those projects are, in the end, unrealized,

inasmuch as that it's their pursuit that matters, it doesn't thereby follow that we might as well abandon them. Abandoning projects is a good way to avoid disappointment, but it is not going to provide anyone with a rich life.

It goes without saying that the permissibility of leg-amputation is insufficient to demonstrate that it will ever be anything like a priority. Resources being what they are, it is possible and likely that the importance of medical interventions like this will always be trumped by more traditionally therapeutic interventions; but that's for another debate.

It's not hard to carry all this back to Bradshaw's case. Her ambition to be a Paralympian may turn out to generate sufficient warrant for the amputation of a disabled, though basically healthy, leg without too much difficulty. (The fact that the leg was disabled makes the decision easier, if anything: after all, some kind of medical intervention was warranted anyway; the question to ask in this case is simply one of how radical we want it to be.) Given that, taking Bradshaw at her word, she would still find her disability a hindrance were she to receive the less radical intervention, it could be argued that the aim of the surgery would be better met with the more radical intervention.

## On me, not in me

So, this argument – if successful – ought at least to give us reason to think that there are times when non-therapeutic biomedical interventions might be acceptable – even if not exactly straightforwardly so – and, more, that they might be justified by making an appeal to the projects that a person might have and around which one might model one's life. Once we've accepted this principle, a way seems to have been opened to all manner of interventions on the body that have the biosciences at their heart.

The *New Scientist* reported in February 2005, for example, that a company called Biojewelry was searching for a couple to have cells from bone fragments grown into wedding rings.[10] Biojewelry was in fact a project devised by Toby Kerridge, of the Royal College of Art, and drew on the medical expertize of Ian Thompson, of Imperial College London. It exploited the fact that it is possible to take living bone tissue and grow it on a lattice, which can be designed according to need. The ability to take bone tissue from the body and grow it according to a deliberate design is obviously of interest from a therapeutic point of view. However, it is easy to see that once you've worked out how to grow bone-tissue *in vitro* according to a design of one kind, the same technique could be used to grow that tissue according to a design of any purpose, well beyond that for which the technology was originally intended. In the Biojewelry project, tissue from the bone fragments would be cultured, and the living cells grown to form

a ring, which would then be decorated with silver. Effectively, this procedure would allow each partner to wear a part of the other: each would be wrapped around the other's finger.

There were objections from the ethics committee overseeing the project that had to be overcome: notably, it was stipulated that if the project were to go ahead, it could only be using bone material recovered during the extraction of a wisdom tooth. (This criterion was not the only ethical point made in relation to the project,[11] but it is the one that is of most interest here.) One can see why such stipulations might be made: recovering bone tissue is an invasive procedure, and it is not entirely surprising that some people should worry about the propriety of invasive procedures carried out for non-therapeutic purposes. Irrespective of whether or not those worries would be powerful, they would be much less prominent in respect of those already facing therapeutic surgery: the risk added to a tooth extraction by removing a sliver of extra bone is likely to be negligible, it does seem to be fairly morally unproblematic for this bone to be removed in the course of a previously indicated procedure. We can present a similar argument in respect of the thought that the removal of the bone sample represents a waste of resources on fripperies when those resources could be put to better medical use.

Hence, when there is some appropriate procedure slated already, concerns about harm will not present any insurmountable objection to taking a little more bone. This does not mean that there is any particular obligation on the surgeon to perform such an extended procedure – he may have no desire to comply, and might point out that the procedure lacks the moral pull of medically necessary surgery. But neither would there be any particular obligation on him *not* to do so that is based on harm avoidance, as there might have been was there no other reason to operate.

And yet this line of thought takes for granted the supposition that the only thing that can serve as a warrant for surgery is medical need. This is a supposition from which I've already distanced myself. Would it be so crazy to consider removing a sliver of bone even in the absence of any medical need?

It is worth looking again at some of the reasons that a person may offer for the justifiability of surgery. Making sense of a person who is seeking surgery involves accepting that he expects that surgery to be worthwhile. Since surgery is unlikely to be pursued for its own sake, its being worthwhile will depend on its being for the sake of something else. This applies just as much to therapeutic as to non-therapeutic surgery, of course; even if the aim of the surgery is nothing more than to forestall death (and whether or not being alive is categorically or hypothetically worth having is beside the point for the moment), it still has to be justified by something beyond itself.

The rings that we'd be growing out of bone might be felt to be valuable for their own sake – worth having in themselves – but it's at least as likely

that they, too, would derive their value from something else. So we need to ascribe to the person undergoing the procedure a belief that the procedure is worthwhile as a means to get bone-grown rings, which are, in turn, worth having because the exchange of bone-grown rings is better in some sense than is the exchange of mere (!) gold ones. The chain of justifications may be very long, and there is no particular reason to suppose that it terminates in anything that is clearly a good-in-itself – *except the pursuit of a particular vision of what a good life would involve*. Yet the warrant for therapeutic surgery also has to be located by tracing a chain or network of reasons for wanting it; and this chain, too, will terminate in a particular vision of what a good life would involve. To this extent, the warrant for surgery for artistic reasons appears to be the same. No obvious reason presents itself for thinking that a sufficient warrant for pursuing surgery in the name of bone-grown jewellery cannot be forthcoming.

Thus we might imagine someone who cheerfully admits that any given procedure – in this case, the kind of thing necessary to harvest bone cells – is risky and injurious, but that this injury is warranted by his project – in this case, the project of obtaining bone-grown jewellery. The argument that his welfare will not be diminished by not having this jewellery will simply bounce off such a person, since his claim was not couched in the language of welfare or wellbeing in the first place. Moreover, though it might be pointed out to him that persevering with the retrieval of tissue could plausibly diminish his welfare, he may be unmoved – again, either because concern for welfare maximization isn't what motivates him, or because he thinks that the welfare loss would be a price worth paying.

Interestingly, we might also present the opposite argument: that that a project involving wellbeing is just as *little* worthwhile as a project involving bone-grown jewellery, inasmuch as that the value of each and any project that a person might adopt sits within a network of beliefs, values and prejudices, none of which is self-sustaining. As it happens, we happen to believe that some projects are important enough to warrant surgery. However, there is no reason to suppose that this belief is beyond criticism or that our values key into anything *objectively* worthwhile.

Of course, I do not imagine that anyone actually *does* suppose that surgery to harvest bone cells has the same value as surgery to remove a tumour. But such a supposition is possible, and it is not obvious to me that it would be based on a mistake or misunderstanding about values. Indeed, it is not clear what the scale is against which we may attempt to measure competing accounts of what is and what is not a worthwhile use of the surgeon's time and art. (This point prefigures a claim that I'll be pursuing in Chapter 8, by the way.) Hence there is no obvious reason why a suitably qualified person might not admit that, while his surgical skills are most frequently put to work in the pursuit of wellbeing – that is, to remove tumours or straighten noses for the sake of

a standard of health, however broadly construed – they might also be put to work in pursuit of some other end – in this case, something artistic. Granted, he might say, surgical intervention *qua* medicine is properly directed to wellbeing – but this does not preclude the possibility that a surgeon may sometimes wear a non-medical hat.

In short, it remains to be seen whether there is any ethical reason why a surgeon might not say that, at least at the moment, he is interested in participating an aesthetic rather than a healthcare project. When we tell surgeons that they are to do no harm without reason, this can generally be taken as an abbreviated way of talking about having sufficient *medical* or *therapeutic* reason; but *generally* does not imply *always*. Almost certainly, a person who argued that surgical harm was warranted because his continued wellbeing was predicated on the possession of bone-grown jewellery would be presenting a weak case. But a surgeon could still take this project seriously without impropriety: to claim that surgeons ought only to concern themselves with medical or therapeutic interventions is indicative, we might say, of an arbitrary prejudice against some projects. (Further, we might add, we risk inconsistency if we are at the same time willing to allow surgeons to perform other non-therapeutic procedures such as circumcision.[12]) Surgery as a means to get jewellery might not be *medically* justified – but this does not mean that it is not justified at all.

## Other objections

The objections that I've considered to harvesting bone cells for the sake of jewellery – and, for that matter, to amputating limbs for the sake of athletic ambition – have so far focused around the idea of whether such operations are particularly sensible from the point of view of the person who desires them. But there're other sources of objection.

For example, one other source of objection to such interventions comes by virtue of a Kass-like appeal to grotesqueness. This is more directly relevant to the case of bone-grown jewellery, but it could be applied to other cases too. One version of the objection would be based on an insistence that there is something grotesque, and so something disreputable, about treating human tissues as the raw material for artworks – or, more generally, that there is something grotesque about treating the body as a raw material for any particular project. Another version would be based on the insistence that there is something grotesque about a part of one person's body being worn as adornment by another.

However, the fact that some practices could be described as grotesque does not make much of a moral difference unless an argument can be produced to tell us, in the first place, that the description is *true*, and, in the second,

that this carries normative weight sufficient to overcome whatever reasons are produced in favour of the allegedly grotesque practice. But since I do not want here to get involved in questions concerning how disputes in matters of taste should be resolved, or whether they can be resolved, or of how one person's essentially aesthetic claims ought to impact on another's behaviour, I shall satisfy myself for present purposes by trying to circumvent such debates altogether by the expedient of adverting to the claim that people who find something grotesque need not be exposed to it if they prefer. The mere fact that something is grotesque is not of great ethical significance.

(Incidentally, if appeals to the grotesquery of using the body as an artistic medium were to turn out to be important, something like bone-grown jewellery would have many potential companions in guilt. The artist Marc Quinn makes a sculpture of his head every five years by casting around 4.5 litres of his own frozen blood – one of these works became the subject of urban legend when the rumour escaped that the sculpture's owner had accidentally unplugged the freezer unit and melted it. Quinn has also made a bust of his baby son, Lucas, from Lucas' placental and umbilical material. Going back a little further, the Italian conceptualist Piero Manzoni sold tinned 'merda d'artista' – artist's shit – as a comment about human production. And sometimes, the body can be seen as a canvas in its own right: all over the world, people tattoo, pierce, and scarify for the sake of adornment. Plainly, at least sometimes, such use of the body is seen as worth doing.)

Maybe the appeal to grotesquery is a slightly unsophisticated mask for something else, though, such that someone who makes such an appeal might turn out to have something else in mind. It might be that, once we unpick it, the appeal is actually an appeal to a worry about objectification, that would run along the lines that we ought not to pursue such things as bone-grown jewellery because to do so objectifies or commodifies the body, and does so problematically. The expanded argument might be that there are serious moral objections to treating persons as commodities; that using bone tissue to make jewellery involves treating the body as a commodity, and so commodifies it; that it is hard to maintain the difference between commodification of the body and commodification of the person; and so we ought not to commodify the body. From the point of view of an investigation into the good life, this kind of worry might bring a concern that a commodified life is unlikely to be flourishing. However, it is not obvious that people who want to use parts of their body in some kind of artistic endeavour really would be commodifying themselves, at least any more than are people who use parts of their body or mind to do their job; it is possible in just about any plausible account of human affairs to treat one's body, labour, mind, talents or whatever as a commodity – something that has a price and can be traded – without thereby treating oneself wholly as a commodity. If people who have to work for a living really are commodifying

themselves, it is not obvious that this is morally problematic, either; we might want to distinguish between the propriety of voluntary commodification and involuntary or non-voluntary commodification, for example. In short, while it may be possible to wrong someone by treating them or their body either as a commodity or as a source of commodities, this does not tell us a great deal at all about the propriety of using the body or tissues in artistic pursuits generally; it simply tells us about the moral qualities of certain kinds of relationship between specific people. (It is noteworthy that, in a passage that seems particularly apposite to the possible use of bone tissue to make wedding rings, even Kant concedes in the *Rechtslehre* that, since marriage is 'the reciprocal giving of one's very person into the possession of the other . . . neither is dehumanised by the bodily use that one makes of the other'.[13]) Since I take it that working is compatible with the good life – particularly if that work facilitates the pursuit of some project of one's own – then there ought to be no particular reason why using one's body in cases like this would clash too horribly with the pursuit of the good life.

Another worry might be that a person could, conceivably, be coerced into undergoing a surgical procedure that they would not otherwise have considered. One of the partners in the Biojewelry project might feel that participation is expected by the other for some reason: participation in the project mightn't be 'authentic'. Such worries may actually be more acute when the thing that will be made is a wedding ring, since there is a particular dimension to the relationship between the person who will wear the piece and the person from whom the raw materials will come that does not apply in many other relationships; we may worry that one partner is, wittingly or unwittingly, pressuring the other into having surgery that he or she would otherwise not choose. In this regard, the ethics committee's concern that the person having the bone harvested ought to require a wisdom tooth removal might actually exacerbate worries. After all, it might well be difficult to persuade an otherwise unwilling partner to submit to surgery if the only reason were provided by the prospect of interesting jewellery – but things could conceivably be a good deal different if surgery was *already* indicated: arm-twisting arguments would seem to be much easier to make if they can be contextualized with a 'since you'll be unconscious anyway . . .'-type appeal.

More generally, there might be any number of other body-modifying procedures that a person could be coerced into undergoing. Once again, though, this objection does not speak to the use of body parts in artistic pursuits in general terms – it only tells us about certain kinds of relationship between certain people. Unless we think – implausibly – that the only way a person might be willing to go under the knife for the sake of some non-therapeutic project is under duress, then we will not have any necessary reason to think that coercion is going to be a big enough ethical worry for us to be able to say that there's something morally problematic about the procedure.

## Body modification and the good life

The argument presented over the last few paragraphs does not, at least on the face of things, speak to questions about enhancement: making jewellery out of bone, or altering the body in any other way for the sake of some artistic project isn't an enhancement in any particularly common sense of the word. Nevertheless, the claims that I've advanced in respect of that kind of procedure are related to claims about enhancement; and they are claims about one of the ways in which the biosciences might allow a person to form and pursue projects of a certain type – and might therefore contribute in some way to securing the good life. And there has been, I hope, a coherent line of argument that's threaded this chapter together.

I argued, in the first place, that there is not necessarily anything sinister about the use of performance enhancers in sport – and, more importantly, that their use is not antagonistic to fairness, the virtues of hard work, or the welfare of the enhancee. I argued that the best justification for enhancing technologies of one sort or another being banned by the governing bodies of a given sport is based on a kind of aesthetic-cum-essentialist argument, the core of which is that there is a set of core characteristics that a given sport has with which enhancement may be in tension. Thus, for example, it may be felt by people who know and care about such things that a certain kind of swimming-costume takes its wearer beyond the limits of competition swimming as it is properly understood.

But this does not mean that those enhancing technologies should not be used at all: it simply means that there may be classifications and subclassifications within a given pursuit that demand a certain set of rules – that there is an 'internal morality' to it. Though some enhancement may violate this internal morality, it does not seem likely to me that many enhancements would: as I've argued, for example, there doesn't have to be a problem with fairness in respect of at least some physiological or pharmacological enhancements. Of those that do, there is still plenty of scope for their use to fall within the internal morality of some other sport: so, for example, if high-tech swimming suits violate the norms of conventional swimming, we might imagine a parallel competition in which they're recognized; if we decide that carbon-fibre blades are not compatible with the particular excellences we want to see in a running race, there could easily be a parallel race for blade-runners – they are not displaying one particular set of excellences does not mean that they display none at all.

But if we recognise something as a subdiscipline in its own right, with its own excellences to admire, then we find ourselves faced with a question: if it is permissible – and admirable – for a person to undergo unusual rigours and hardships for the sake of achieving glory in one sporting arena, is there any coherent reason why it would not be permissible or admirable to undergo

unusual rigours and hardships for the sake of achieving glory in another? If a person adopts this particular kind of life as being important to them, and assuming that they have not made a mistake in doing so, then there is at least a *prima facie* reason to think that it would be good for them to pursue the means necessary. Since the pursuit of projects is at least a component of the good life, my claim is that not only would such pursuit be permissible – it might be the sort of thing that is necessary for them to flourish according to the standards they have determined for themselves. And, using the case of Danielle Bradshaw, and an imaginary athlete called Paddy, I've suggested that this kind of thought applies even if that means a fairly radical remodelling of the body. For sure, the more radical and the less reversible the procedure, the more tempting it may be to worry about its wisdom – but there is nothing in principle objectionable to it.

It does not quite follow from a claim that $\phi$-ing is permissible that helping someone to $\phi$ is permissible as well – there is a logical difference – but it seems to me to be only a small moral leap that has to be made to say that it is therefore permissible for a surgeon to undertake to remove even a healthy and functional pair of legs for the sake of helping a person realize her project of a particular sort of sporting glory. That is to say: a person's project may warrant some fairly extreme bodily interventions, some of which are of the sort that are only possible because of, and require, bioscientific expertize of one sort or another. There need be nothing particularly ethically problematic about this; and there need be no antagonism between such procedures and the good life. Once this principle is accepted, there would seem to be very little reason to deny that it could be extended to cover things like bone-grown jewellery – or any other pursuit that requires the modification of the body and/or the manipulation of body tissue for the sake of the realization of a project. The constraint that applies is whether that project really is contributory to the good life, and whether the means chosen to pursue it are particularly wise ones.

But the important point is that the biosciences do offer us the means to realize any number of (admittedly potentially weird and wonderful) projects; and there need be no particular worry about making use of the biosciences in such a way unless one adheres to the idea that the proper use of the biosciences is therapeutic. Since the pursuit of projects is a part of the good life – and, more importantly, since there is no reason to suppose that the non-therapeutic use of the biosciences is a misuse – we can say that the biosciences do promise to contribute to the good life in a manner beyond their therapeutic potential.

In other words, prolonging a person's life indefinitely may not be a staggeringly good use for the biosciences; but making a bodily adornment may be better.

## Notes

1 Austin, 2009, *passim*.
2 Ibid., p. 50.
3 Though this is not to deny that a motor may, of course, give an unfair advantage as well. In this sort of context, it's interesting to consider the use of carbon-fibre blades by amputee athletes like Oscar Pistorius; inasmuch as those blades will not tire in the way that muscle tissue in even the most well-trained 'normal' athlete will, it is not implausible to suggest that, *ceteris paribus*, they will generate an advantage; whether or not it's unfair may be a moot point: the IOC, in allowing him to compete alongside and against normal runners, appears at the moment to think not.
4 Nussbaum, 2006, pp. 180–2.
5 Sandel, 2007, p. 30.
6 Anon, 2009a; emphasis mine.
7 Anon, 2010a; slightly adapted.
8 McKeegan, 2010.
9 For a discussion of some of the ethical issues related to scarification, see Oultram, 2009, *passim*.
10 Horgan, 2005, p. 15.
11 Personal communication.
12 On which point, see British Medical Association, 2004.
13 Kant, 1998, 4:359.

# 6

# Thinking Better about Better Thinking

In this chapter, I shall look at the possibility of cognitive enhancement and the contribution that it might make to flourishing. Cognitive enhancement is, in the words of Nick Bostrom and Anders Sandberg, 'the amplification or extension of core capacities of the mind through improvement or augmentation of internal or external information processing systems'.[1] Their understanding of cognition extends to methods of acquiring information, processing it, remembering it, and putting it to use; as such, '[i]nterventions to improve cognitive function may be directed at any one of these core faculties'.[2] For the sake of this chapter, I'm going to accept their definition of cognitive enhancement, though I will avoid talking about new ways to acquire information – though Bostrom and Sandberg do talk about direct brain-computer interfaces, they do so in the context of using the brain directly to do things in the world[3]; and though we could imagine one day there being technology that would reverse this direction of flow and allow us to upload new knowledge and skills directly, it is too speculative at the moment to warrant much attention here.

From the off, I shall take as read that there are situations in which altering a person's mental processes is desirable. Notably, there are conditions such as Post-Traumatic Stress Disorder (PTSD) in which a person's mental processing of certain events is disrupted, causing their memory to intrude painfully into present life. There is some evidence that it may be possible to ameliorate this condition by the use of drugs such as MDMA (ecstasy), at least as an adjunct to psychotherapy; thus Mithoefer et al. state that

> there [is] a favorable long-term risk/benefit ratio for PTSD treatment with just a few doses of pure MDMA administered in a supportive setting, in conjunction with psychotherapy. Should further research validate our initial findings, we predict that MDMA-assisted psychotherapy will become an important treatment option for this very challenging clinical and public health problem.[4]

Such interventions should be welcomed as ways to help people put right an aspect of their life that had gone wrong. Relatedly, drugs and implants might be able to slow or even halt the progress of age-related cognitive decline. However, such interventions should be remembered for what they are: attempts to make an aspect of a life that is currently not acceptably good better. As such, they too fall clearly on the 'therapy' side of the T/E distinction, and so their desirability won't tell us a great deal about cognitive enhancement *qua* enhancement. What

remains to be seen is whether such interventions would provide all that much benefit to people whose cognitive capacities *are* at present at least acceptably good: whether they can help maximize flourishing by increasing the cognitive capacity generally.

The hypothesis that I'll be pursuing here is that their role will be fairly minimal. There's a couple of aspects to the argument. First, and least powerful, is that deciding what actually counts as a cognitive enhancement is not straightforward. Particularly in respect of memory, what it takes to make a person's life better is up for debate. But, more importantly, I'll be suggesting that, in many cases, bioscientific cognitive enhancement does not offer anything that could not be got non-bioscientifically – and what it promises may actually be parasitic on those non-bioscientific resources. Furthermore, in at least some of the circumstances in which cognitive enhancement may make a difference to someone's life, we might have independent reasons to doubt that it's necessarily the right strategy to adopt.

A hint about why I think that we should not get too excited about bioscientific cognitive enhancers comes from the fact that we can legitimately call a great many everyday devices cognitive enhancers. A pocket calculator is one such: it helps us perform fairly complicated arithmetical functions more quickly and more reliably than might otherwise be possible, and its memory function saves us the burden of remembering variables or formulae. More simply, a calendar, either paper or electronic, might also count as a non-intellectual cognitive enhancer, inasmuch as that it helps us organize our day. Even writing itself can be seen as cognitively enhancing: it allows us to keep a record of our own and others' thoughts, to communicate at large distances and over time, and so on, and allows us to solve the problems presented by a fallible human memory: it is much easier to remember the 20 or 30 characters of most alphabets than to try to remember the things that they may be manipulated to represent. People whose language uses non-alphabetic script have a harder time remembering what sign means what, and may not know how to decipher a word they've not encountered before – but it remains true that *any* kind of system that can represent an idea graphically dramatically reduces the mental burden of those who use it. Moreover, a character once learnt can be used in any number of contexts; and so even if your script requires that you learn many characters, that will still take less effort than having to learn by rote the content of every document in which they may otherwise have been put to use. Famously, it was precisely its role as a substitute for memory 'proper' that made Plato suspicious of writing.[5] Given that we are familiar with a range of cognitively enhancing technologies, then, the question is this: is there any need, or indeed any benefit, from pursuing enhancers of a different kind by manipulating the brain and its processes?

It's with the idea of improving our memory that I'll start; I'll argue that there really isn't much to the idea that our lives would be much improved by means of bioscientific mnemonic enhancement.

## Enhancing memory

Forgetfulness is annoying, even when it's not an indication of anything particularly medically sinister. It is not hard to come up with reasons to think that having our memories improved would represent a good and desirable use of bioscientific interventions, and technologies that could make us better at remembering things seem straightforwardly good, even in those for whom such interventions would not count as therapeutic; there is a level to which we are prepared to tolerate lapses of memory, but this doesn't commit us to thinking that it wouldn't be a benefit not to have to tolerate them.

Offhand, the benefits of mnemonic enhancement would seem to be considerable. On a workaday level, our lives would be made better – not necessarily by much, but by a noticeable amount – if we were less likely to struggle with passwords, vague acquaintances' names, and so on. Farah et al. point out that

> [t]he ability to encode new memories declines measurably from the third decade of life onwards, and by the fourth decade the decline can become noticeable and bothersome to normal healthy individuals. Memory difficulties in middle or old age are not necessarily a harbinger of future dementia but can be part of the normal pattern of cognitive ageing, which does not make it any less inconvenient when we misplace our glasses or forget the name of a recent acquaintance.[6]

On a less workaday level, too, we might find significant benefits. For example, it is not hard to see how it might be highly desirable to have a better recall of something like the grammar and vocabulary of a foreign language; and there are times when we are better off for being able to bring to mind certain facts.

But, once we begin to scratch at the question a little, is there any guarantee that improvements in memory would make that much difference to our lives? Possibly not.

An important thing to keep in mind is that there is more to memory than the mere retrieval of detail. Memory is important at least significantly (if not wholly) because of what it does for us and the work to which its put. If memory matters to us, then implicitly it's because we want it to reflect reality. But this raises epistemic issues that open a range of questions about whether mnemonic enhancement would be able to do all that much for us. Someone who wants his memory to be improved is, implicitly, not interested simply in there being a narrative about things that happened in the past; it matters to him that the narrative is reliable. He should be able to tell the difference between a memory and a pseudo-memory: should it turn out that one of the things he 'remembers' didn't happen quite that way, he'd be disappointed. (Were he not to be disappointed, then there would seem to be that much less of a reason to seek mnemonic enhancement; confabulation would do the trick.) But this, at least sometimes, might be a real issue.

The worry here is suggested by some of Wittgenstein's comments on the impossibility of a private language. The idea here goes something like this: one can only really be said to understand a language when one can tell when the words within it are being used correctly or incorrectly. But a person who attempts to use a notionally private language – a language that only he speaks – would not be able to do this. Imagine, for example, that someone decided to formulate a private language and set about the task of assigning words to name the things in his world; call him Adam. Adam looks around the room and decides that *this* will henceforth be called a *zolog*, *this* an *urpisi*, those things *tirizmef*, and so on. He then decides – reasonably enough – that he ought to check that his new vocabulary has sunk in; and so he starts looking around the room again, repeating to himself the names he has invented. But how can he check that his recap matches the words he invented a little while ago? By what standard can he be certain that the object he is now inclined to call an *urpisi* isn't what he called *zolog* before? And – come to think of it – was it *zolog* or *solog*, anyway? Was *tirizmef* singular or plural? Was it even a noun or an adjective?

Adam can't check his use of the language against his memory of it, because it's his memory of it that he isn't sure of to begin with. (Trying to check that he has remembered his language aright by comparing it with his memory would be a bit like, to echo Wittgenstein, buying a second copy of the morning newspaper to make sure that the first reported the facts accurately.[7]) And since he can't in all honesty say that he can tell whether or not he is using even his own vocabulary correctly from one moment to the next, and has nothing against which to check it except his using it, he can't claim to be able to know his 'own' language.

Note that what's important here is that it doesn't obviously matter whether, at a matter of fact, Adam's recall is accurate. What matters is that he would never have any particular grounds to *think* it accurate if he remained the only source of authority about the language. In order to be able to say that he understands it, he has to be able to cross-refer his use with something outside his own head. There needs to be an external point of reference. In respect of languages, this reference point might be found in the everyday use of a word, or in something a bit more formal: we might imagine that, had Adam planned ahead, he would have used a notebook to record in English the words and grammar of his new language.

What applies to Adam and his use of language seems to me to apply to someone who has had her memory 'enhanced'. Let's imagine that Adam's partner – whom I'll call, inevitably, Eve – decides that she will make use of a new drug that promises to make her memory much more powerful. It certainly feels to her that details about names, bank PINs, and so on come to mind much more easily. But Eve, having been laughing at Adam's attempts to teach himself his new language for some time now, recognizes that those details could be

wrong: she might just *think* that she remembers them. Like Adam, what she really needs is some authority outside her own head to show her that what she *thinks* is a memory is not a pseudomemory after all. The point is that, if there is some fact about which we might conceivably feel a more-than-background level of uncertainty – and that covers all facts – it would do us well to have some extra-cranial standard against which we could check it.

## Out of our heads

Now comes the important part: if people or things outside of a person's head are necessary to serve as the knife with which we can pare memory from pseudomemory, we can say that that outside world is (in a sense) a part of memory; the external reference corroborates the internal archive, thereby helping to ensure that the passing show of mental images, phrases, and so on really is a memory properly so-called. And each of these claims has important implications when it comes to assessing the point and benefit of mnemonic enhancement. After all: the external reference point is constituted by things like Adam's lexicon and Eve's downloaded timetable. It could also be found in other people: just as we are kept 'on track' in our linguistic habits by the network of communication of which we are a part, the same can apply to our memories. And these reference points would, quite obviously, be unaffected by mnemonic enhancers. Though such enhancers may mean that our 'internal' memories are, as a matter of fact, more likely to be reliable, the important point is that we will only be able to *say* that they're reliable with any confidence thanks to the possibility of checking them against the external world. And this constitutes, at the very best, a limit on the benefits that are likely to be derived from mnemonic enhancement.

Second – and following on from this – the outside world can be treated as an archive in its own right. A person does not have to remember everything in order to enjoy the benefits of a good memory: he just has to have access to a mnemonic repository of some sort. Just as the technology of writing means that one can have access to all the world's information without having to remember anything more than a few characters and some punctuation, so there are plenty of comparable 'extra cranial' archives. Other people can also figure as constituent parts of this archive, too: they can serve as means to fill in things that we can't remember ourselves, as well as providing a way to corroborate or dispute memories. What matters is that, if some aspect of the publicly-accessible world is required as a means to corroborate or disprove certain memories, those aspects themselves can be treated as stores of information. If there's something that we can't, but want to, remember, the simple thing to do is to ask someone who might be able to help fill in the gaps.

And this must prompt a question: can we really be so sure that mnemonic enhancements of a biomedical sort would provide us with all that much of a benefit? Without comparing our enhanced internal archive with the world around us, we won't be certain that it's any good as an archive; and given that the contents of any internal archive can be 'contracted out' to an external archive of books, electronic devices and friends, it's not at all clear that there'd be all that much to be gained from an enlarged internal archive to begin with.

What can't be stored in this way are the 'qualia' of memory: not the fact that your personal hero died on this day, but what it felt like to hear of the death of your personal hero. But this kind of memory couldn't be stored with any degree of reliability anyway: only the report that you felt that way. And if you do happen to fool yourself about those feelings: well, who really cares? Perhaps only those with a heavy investment in the right kind of concept of authenticity – but since memories wither when they are kept private anyway, this investment may only ever be dubiously well-advised.

Admittedly, there is likely to be a number of shortcomings with the external archive. For example, it might be that none of the witnesses to an event has a particularly good memory of what went on, and that the account we reconstruct is not true. The external archive is not static: just as words change their meanings or acquire new connotations, so 'crowd-sourced' memories may move around as well. Often, this won't matter – but sometimes it might: if the police are trying to reconstruct what happened at the scene of the crime, then a number of slightly unreliable accounts from witnesses might, collectively, generate a seriously misleading account. I'll return to this problem in a moment – though, for the time being, it's enough to note that this limitation does not directly link to any claims that we might make about the good life.

At other times, the shortcomings *do* have more direct impact: a student sitting an exam frequently can't consult with books and invariably can't consult with the person sitting next to her; a novice barman will have a need to learn the prices of the drinks he sells as quickly as possible. Hence we might think that, at least sometimes, it would be desirable to boost our own mnemonic capabilities to maximize our independence from the external archive.

This seems fair enough; but none of the examples above will show that mnemonic enhancement would be a particularly life-improving thing for most of us most of the time: our fallible memories are for the most part good enough, and when they give out, we can turn to the external archive. More importantly for the people in these examples, though, it is not obvious that the biosciences are necessary to make lives better: the barman's life can be improved, for example, by means of an electronic till that remembers prices for him, and then adds them up more reliably as well – even when the bar is noisy and busy. The most difficult confounder is presented by the student's desire to cram in a little more information. Still, this isn't enough to dent the general point; and, of course, we may want to question the educational value of

recall anyway: it's unlikely that there will ever be another time when a lawyer or doctor would not be able to use reference tools; what seems much more important is an ability to locate and use them. Remembering how to use a library or search engine doesn't require enhancement, though: rather, someone who can't remember that requires therapeutic intervention; and so her case doesn't really speak to the concerns I have here.

## Criminal detection: A duty to remember?

Things are a little more complicated when we consider the possibility of mnemonic enhancement in the context of criminal detection. It might be that, if people had more acute memories for detail, it would be easier to reconstruct a chain of events, and as a result, it might be that justice is served slightly more effectively. (Wagenaar is sceptical about the possibility of such enhancements[8]: but we can still consider whether they'd be desirable should they be possible.) This may not be wholly desirable, though – at least on the grounds of a concern to improve the life of the enhancee. There is a couple of reasons for this. First, the witness to a traumatic scene might not benefit from being able to remember the details: it might be highly desirable for those memories to drift off and be diluted by time's passing, much as applies to any other memory. We've already noted that traumatic events that retain their currency may lead to PTSD and a need for therapeutic intervention. ('Mnemonic enhancement' might even amount to a heightened ability to forget – and here it's tempting to think of the film *Eternal Sunshine of the Spotless Mind*.) Second, even if there is no particular trauma associated with the memory, we would still be thinking about situations in which memory was enhanced not for the sake of the enhancee, but for the sake of some other end.

On the other hand, if flourishing is to be measured in terms of the capacity to perform typically human functions and roles, and to live better is to perform them better, this does open the way for a counterargument. For among human functions and roles is the ability to identify and discharge a duty. And it might be that we have a duty to remember for the sake of criminal justice. If it's plausible that the good life is a life in which one performs one's office, we might have a duty – or at least a reason – to remember after all *as a part of that life*. This is not a small consideration. Neither is it a small consideration that a functioning criminal justice system might be part of a well-ordered society, and membership of a well-ordered society might be a criterion of human flourishing; this would give an indirect reason to ensure that memories were as accurate as possible. But the point remains that memory in this context would be only a part of human function, and only a part of human flourishing. There is, for sure, room for debate about the strength of a 'duty to remember'. The point for the moment, though, is simply that it's a debate that must be had.[9] We can't take anything off the shelf.

## Memory and absentmindedness

The position at the moment is this: that there is room to be doubtful about whether neurological or pharmacological interventions designed to increase mnemonic capacity in individuals are in general really worth the bother. We might very well be able to augment their recall by making use of external resources, from diaries to conversations to photographs; but though such technologies do fit into Bostrom and Sandberg's 'amplification or extension of core capacities of the mind through improvement or augmentation of internal or external information processing systems', they seem hardly to be anything worth getting excited about. And it's not clear that we should get particularly excited about more high-tech ways of helping us not to forget, either. They don't seem to be so radically different in kind.

Before leaving mnemonic enhancement to one side, it's worth noting that recall is not the only thing that brains do, and the process of forming memories is not simply a matter of layering facts one on top of the other. We can also lose track of what we're doing, and thereby fail to form or consolidate memories to begin with. For example, if it's a warm day, and he's a lot on his mind, and his keys and wallet are in his trouser pocket, it might well be very easy for a person to leave a jacket behind somewhere and not notice for days afterwards – by which time, it's gone for good. The mnemonic failure here is not so much one of accurate recall; it could well be one of processing. For this reason, it's tempting to think that, if we could increase the brain's 'handling capacity' – its ability to keep several plates spinning – we could expect to be less forgetful. And, as such, it's tempting to think that making use of some kind of enhancement that could stimulate the brain to be able to keep track of more tasks at once would be highly desirable and life-improving.

However, it might be worth reining in our enthusiasm a little bit even here. The person who leaves his jacket on the park bench presumably has things playing on his mind that are, at least at that moment, more important; had there been less competition for his attention, he would have remembered it. So far so good.

A little extra mental capacity might have meant that comparatively small things like the jacket would not have slipped his attention. But is it obvious that dealing with a much greater number of things all at once really would be a trait worth cultivating? There is certainly an annoyance factor that might be avoided. However, it does not follow that never losing track of things is all-things-considered a benefit. One might speculate that, notwithstanding the occasions on which we trip up and we wish it hadn't, it could be a good thing that our brains have the ability to prioritize some things, and to push others aside: without it, we would find ourselves overwhelmed by calls on our attention, and end up intellectually paralysed.

In the story 'Funes, the Memorious', Borges famously and playfully cast doubt on the supposition that more memory is good; since he could not put to one side particular traits, it was difficult for Ireneo Funes 'to understand that the generic term *dog* embraced so many unlike specimens of differing sizes and differing forms'.[10] But one does not have to advert to fiction to question the idea that more memory is desirable, even when it is not traumatic. Citing the work of Daniel Schacter, Viktor Meyer-Schönberger suggests that

> [u]sing generalizations, relying on conjecture, emphasizing the present, and respecting subsequent experiences, helps us to reason swiftly and economically, to abstract and generalize, and to act in time, rather than remain caught up in conflicting recollections.[11]

It is admittedly possible that Meyer-Schönberger is romanticizing a little. But even given that possibility, there might be people who do have too many recollections.

For example, 'AJ', the subject of a paper in the journal *Neurocase*, can remember in astonishing detail facts about her biography from the age of 11, and is described as 'both the warden and the prisoner of her memories'.[12] She reports that 'I can't let go of things because of my memory' and that '[m]ost have called it a gift but I call it a burden. I run my entire life through my head every day and it drives me crazy!!! . . .'[13]. Her memory does not provide her with any particular compensating benefit, though: oddly enough, she even has difficulty with some aspects of recollection:

> AJ reports that she never excelled in school and says she "hated" it. She reports she had trouble memorizing dates in history, arithmetic, foreign languages, sciences and "got Ds in geometry." Memorizing poetry was painfully difficult. Yet she can recall with ease every one of her teachers since kindergarten.
>
> In contrast to AJ's strong autobiographical memory, and her ability to recall dates and events, she is not a gifted memorizer. For example, she told us she has five keys on her keyring and can never recall which key is for what.[14]

So – absent the suggestion that there is anything pathological about the things AJ can't remember – an improved recall is not necessarily indicative of better memory in another sense; it may even be undesirable, and so it wouldn't follow that we could always make our lives better by expanding our mnemonic capacity.

In effect, if the possibility of annoying absent-mindedness is the price we pay for having the ability to ignore unimportant things – which is, presumably, not only a desirable ability, but itself one that we might want to enhance in some cases – then it might be that it's a price worth paying (even if that does mean having to buy a new jacket once in a while). Neither – to return to the jacket example – would it be as simple as all that to say that what would be really good is a heightened ability to prioritize better, so that we don't relegate things that will later turn out to be important – not least because we often don't have

that kind of predictive power, but also because emphasizing what might be important in the medium-to-long term might mean getting impossibly bogged down in the short.

Slightly more speculatively, we might wonder whether we can get carried away with the idea of mnemonic accuracy as the benchmark of desirability, and whether doing so would be to make a mistake about the role played by the memory in a flourishing human life to begin with. For sure, information retrieval is something that we want, and an ability to retrieve information that falls below a standard of reasonable acceptability is a problem that we might do well to solve. But humans are not information retrieval systems, and the fact that we have built machines specifically to be better at retrieving information than we does not tell us that it would be good to aspire to match those machines. Indeed, the fact that we do have such systems speaks to the relative unimportance of working to extend the range of the capacity of the human brain: what is the need to remember things given the ubiquity and ease of data-retrieval systems far more capacious and reliable than even the most boosted human could ever be?

More positively, people have a character that a PC simply cannot match; and a slightly 'lumpy' memory is likely to be one of the things out of which that character is constructed. Importantly, a lumpy memory allows for serendipity – not only is it the case that overhearing a snatch of a conversation may, for no particularly sensible reasons, help us home in on a thought that's been nagging away for days, and this would not be possible were our memories more 'efficiently' stored; it might also be the case that being unable to put our finger on quite what it is that we want leads us down intellectual alleys that may otherwise have remained unexplored. In other words, lapses of processing that manifest themselves as lapses of memory may be the spur to creativity. (I'll return to this point in a little while.) Making memory more efficient may turn out to be no enhancement at all.

## Enhancing processing

There is, of course, more to cognitive enhancement than remembering things. At least as important is the ability to *do* things with the information to which one has access: in the bluntest terms, cleverness counts. Should it turn out that there is a way to increase processing ability, it's easy enough to assume that it would be of benefit to make use of it. There's a range of considerations that might be brought to the table to support this assumption. Some relate to the agent directly – that is, they relate to benefits that we might think accrue simply by virtue of increased cleverness. Some relate to the agent more indirectly – that is, they relate to the possibility that though cognitive enhancement *per se* may not make an appreciable difference to the agent's quality of life, it does generate other opportunities that will turn out to benefit him.

One reason that we might have to pursue cognitive enhancement would be that it would give us a positional advantage over others in, say, the job market, and so contribute to the possibility of being more comfortably off in the long term. We might worry that emphasizing the positional benefits of cognitive enhancement could possibly erode other good things, such as solidarity. But a person who makes use of a cleverness-boosting procedure is not likely to suffer on this count; while someone at the bottom of the intellectual league may be isolated, the uniquely intellectually capable are not, I would contend, in anything like that position. Indeed, cognitive enhancement may mean that at least some people who would otherwise be alienated from the community around them are able to play a much more active and equal role in it, and so a barrier to their flourishing *qua* social beings may well be removed. Being better than your rivals is perhaps desirable; but being as good as them is also worth pursuing if you wouldn't otherwise be.

But not all the benefits of cognitive enhancement that accrue to the enhancee need be positional. It's also tempting to think that cognitive ability is its own reward.

The thought here is that a person is simply better off by virtue of having a deeper cognitive well from which to draw and that it is better by some standard to be more cognitively able than less. This kind of thinking is not entirely without intuitive merit, and at least on the face of things, it would not be difficult to generate a tolerably robust argument in its defence on a range of grounds and be in good company. Mill thinks that it is better to be Socrates dissatisfied than a fool satisfied.[15] Were Socrates a fool, it would be hard to see why it would be better to be him than the other fool – that would appear to amount to the idea that it's *ceteris paribus* better to be dissatisfied than satisfied, which is a little too strange to accept. So we have to assume that Mill is thinking about the famous Socrates, who is not a fool. In that case, it's tempting to infer that Socrates' greater wisdom is the kind of thing the value of which can offset his dissatisfaction simply by being there. In a more Aristotelian mood, we might argue that, since intellect is one of the defining characteristics of humanity, it is therefore simply a good thing to have *qua* human.

Yet, although I would not want to dismiss this sort of consideration out of hand, I think that it is worth being a little circumspect. Even if we think that the benefits of intellectual ability are measurable in non-positional terms, it might be that each additional 'unit' of enhancement contributes less to the flourishing of the person who is enhanced. This being the case, the reason to get further enhanced would diminish with each enhancement, until at some point the possibility of extra-cognitive ability would be a matter of indifference to the would-be enhancee. This might make a difference to our decisions about whether there is much point pursuing the enhancement. (If Socrates is dissatisfied not despite, but because of his wisdom – maybe his knowledge that he knows nothing is the source; the fool is even more restricted, for sure, but

his ignorance is one of the things of which he is ignorant – it might even be the case that extra brainpower might have a marginal malefit!)

A case for there being a diminishing marginal return on extra processing capacity is fairly straightforwardly made. While there may be people who are cognitively retarded in such a way as to mean that they struggle with more or less every aspect of everyday life, and in whose case some intervention that increases cognitive ability may plausibly make a significant difference to flourishing, providing an intervention that is capable of generating the same improvement in cognitive ability to someone who is able to cope perfectly well with everyday intellectual demands may, with equal plausibility, make much less of a difference to their life; after all, if the world demands no more of them, their extra ability may never generate any extra benefit. Having extra cognitive capacity in no more challenging a world may be like having a million extra pounds when the shops are shut. Now, the retort to this is that, in a world that is ever-more cognitively demanding, even the highly-functioning may well benefit from *some* enhancement. I'll return to this retort towards the end of the chapter: such a world may itself be antagonistic to flourishing.

For the time being, it's enough to raise the question of whether being more intellectually able would always generate benefits. It's true, of course, that there's a range of things that I might be able to do, or do better, if I could boost certain aspects of my intellectual capacity: cryptic crosswords have always been a weakness, and they generally leave me baffled. Maybe that's just because I don't attempt enough, and I would improve with practice; but let's imagine that my failure is attributable to some part of my intellect that is stuck in a low gear. Maybe, on the basis that it's better to be able to do things than not, my life would be improved by being able to shift this part of my mind up a few gears. But it is, I admit, hard to see how this would really improve my life. Perhaps the problem is, as Mill indicates in respect of the fool and Socrates that I simply can't see what I'm missing.[16] In that case, however, I am not in a position currently to say that I'm missing anything. Were there to arise some identifiable need to complete cryptic crosswords, I might have more of a cause to desire enhancement in that sense. At the moment, I don't; the same applies to other tasks that cleverer people might be able to perform. But Socrates is dissatisfied, and I have yet to be convinced that I would be better off if I were too.

Even in respect of tasks that we want to perform but for which we do not have the cognitive wherewithal, it is not obvious that enhancing by tweaking the mind would have much of a role to play in realizing our desires. Just as things like writing and the microchip allow us to outsource our mnemonic capacities, so there is a range of ways in which we can boost our processing abilities externally. Adding long lists of numbers is easier with a paper and pencil; more complex tasks – say, working out the likely path of a hurricane – require a computer. Nor is it obvious that there need be any less satisfaction derived from being able to get these 'external brains' to work for us. A physicist

who has worked out the mass of a new particle may have been unable to do the work without sophisticated electronics; but that is hardly likely to diminish his sense of achievement, or rule him out of the Nobel prize. (There's a shade of the Real Tomato Argument mentioned briefly in the last chapter here: the achievement is undiminished, even if it did come with assistance.) Once again, the hypothesis that suggests itself is that cognitive enhancement is not necessarily the way to increase the goodness of a person's life; and – perhaps more importantly – that even if an enhancement is in order, it isn't obviously going to have to be of the sort provided by the biosciences.

One other sceptical worry about the claim that cognitive enhancement benefits the individual is worth mentioning, and it's to do with whether there's a pressure to enhance. Quite obviously, if we make use of cognitive enhancers for the sake of generating positional benefits, we have to be exceptionalist about things: we have to assume that we are the only ones using the technique (or any comparable technique), lest the positional difference be lost. But the converse of this is that, if there's a reason to believe that others *are* using enhancers, it will generate a pressure to enhance for the sake of not lagging behind. Is this necessarily a bad thing overall? Well, it might be, if a person – call him Bob – who was otherwise perfectly satisfied with his level of intellectual function felt pressured to enhance himself when he would otherwise not. The counter to this is to echo Harris' point about educational enhancement that I mentioned in Chapter 3: if being clever is better than not, the non-positional benefit would still be there for Bob to enjoy. But this does perhaps miss the point that, if Bob was not motivated to enhance himself cognitively by the non-positional advantages that doing so offered, these same advantages would not obviously strike him as much of a compensation for having been pressured to do so now. For sure, enhancement in Bob's case would have meant that he would not lag behind – but that isn't compellingly going to be the sort of thing that adds to flourishing so much as it's prophylaxis against the possibility that one might flourish less due to others' actions thanks to positional loss, erosion of solidarity, and so on. If being pressured to enhance erodes Bob's sense of freedom, or his sense of the kind of person that he is, it might even erode important aspects of his flourishing.

## The argument from alienation

One of the arguments sometimes put forward in favour of cognitive enhancement is that it will maximize people's ability to live and flourish within an increasingly cognitively complex world. Quite obviously, modernity's characteristics are attributable to the workings of the human brain; but that brain itself evolved gradually, over thousands of years in which not very much happened that would demand significant high-order thought – and so the idea that the human brain is ill-suited to the demands of the world today is not

wholly surprising, and nor is it wholly implausible. Neither is it surprising that there should be suggestions that we could use cognitive enhancers to allow the brain to function in a manner more attuned to its environment.

In this vein, Bostrom and Sandberg write that

> if modern society requires much more study and intellectual concentration than was typical for the human species in its environment of evolutionary adaptation, then it is unsurprising that many people today struggle to meet the demands of the school or the workplace.[17]

And if people really do struggle to meet the demands of the school or the workplace, and do so for the simple reason that they are human and human brains have been shaped more by thousands of centuries as hunters and pastoralists than by a few centuries as scribes, engineers and insurance brokers, then it is perhaps fairly natural to wonder, in the first place, whether humans today really are flourishing as much as they might, and, in the second, whether their flourishing might be increased by giving their brains a helping hand by means of some biomedical intervention. 'Technological self-modification and the use of cognitive enhancement methods', Bostrom and Sandberg continue, 'can be seen as an extension of the human species' ability to adapt to its environment'.[18]

This optimistic statement is not without its problems.

One slightly naïve response to Bostrom and Sandberg's optimism might involve wondering aloud whether giving people enhancements to meet today's requirements might yet turn out to be a bit of a waste of time in respect of tomorrow's. A person will derive no benefit from being better able to flourish in environment *e* if the society in which he has grown to maturity now looks on the characteristics of *e* as being quaint and obsolete. However, this line of objection relies on having placed the weight of the argument on specific capacities; Bostrom and Sandberg are considering enhancements *in abstracto*, and we can easily imagine enhancements that have a high 'teleological elasticity' – that is, enhancements that would give the recipient a higher performance capacity whatever the situation in which they were put to work.

However, the optimism of Bostrom and Sandberg's account does suggest that it is always the best thing to alter people to fit their environment; and this is a supposition that invites at least a degree of scepticism. Very plausibly, there are situations in which we might well think that altering people to meet the demands of the environment is not the way to go if we really are concerned with the good life. The reason for this is a concern about alienation. By this, I mean a situation in which people are enhanced as a means of achieving a good life understood not so much in terms of what a human life is or could be, but in terms of what it has to be given the socio-economic context.

Imagine that we lived in a world in which most humans were expected to work 100-hour weeks in hot and cramped conditions. Call this world 'Metropolis'. It is generally recognized that people do not flourish under such conditions; but a scientist (call him Rotwang) offers to make available an additive he has developed which, when added to the water supply, will subtly alter the neurophysiology of the people who have consumed it, so as to mean that they are better able to cope with the environment in which they find themselves. Would the flourishing of the people of Metropolis be enhanced by the use of Rotwang's additive?

It is tempting to say that it would be. It is better not to be exhausted and ground down by one's environment, and anything that can help avoid this ought to be seen as assisting people's flourishing. However, we should perhaps be a little wary of helping ourselves to this conclusion without further thought. For it assumes that it is desirable to keep the prevailing socio-economic conditions in Metropolis as they are, and that improving the lives of its citizens is a matter of fitting them better around those conditions. From the point of view of the good life, this is far from being self-evidently desirable strategy: for example, we could imagine someone advancing a claim that the good life properly understood involves and requires negative attributes like not being ground down by one's environment, but also positive attributes like the cultivation of intellectual and aesthetic capacities. Even if the negative requirement is satisfied by Rotwang's additive, it would appear tempting to wonder whether they would be living a particularly good life in the absence of the positive. Such an argument appears not to be wholly without merit.

More generally, a world in which we are prepared to make use of something like Rotwang's additive rather than alter the socio-economic conditions looks like a world in which there is an elevated chance of people being made subordinate to those conditions; and there are several arguments available that might lead us to suppose that that, itself, is antagonistic to flourishing. It's in the light of this kind of consideration that we can tell a story about alienation – the word is deliberately chosen because of its Marxist echoes[19] – inasmuch as that the shape of the enhancee's life is determined by others' (perhaps an impersonal other's) models. Note this is different from the picture outlined in Chapter 1, in which a person might be persuaded to revise his account of the good life by means of public deliberation, because the worry here is that there *is* no deliberation. We could tell a similar version of the same story framed in Kantian terms, expressing a concern that Metropolis fails to treat its people as ends in themselves; or we might be attracted to an account inspired by Heideggerian worries about treating people as 'standing reserves' in the face of an overweening economic system.[20] Each of these accounts has its problems; but the point is that it is not unreasonable to wonder whether, in the first place, there are certain socio-economic systems that are antagonistic to flourishing (at least in any rich sense) and, in the

second, whether a system that is willing to make use of Rotwang's additive might be one of them.

The point is that simply talking about people adapting to their environment is not, by some way, the whole story. We need also to consider whether the environment is one to which we would reasonably want to adapt. Now: is today's world one to which it would be reasonable to want to adopt human capacities?

The world described by Bostrom and Sandberg, even though many people may struggle to meet the demands of the school or workplace, is probably not like Metropolis: the cognitive demands placed on us are well within the capacities of most people; and this, in turn, counts as some reassurance that a failure to cope with the world's cognitive demands is properly thought of as requiring therapy rather than a cause for enhancement as I outlined the distinction in Chapter 1. There is a case to be made for providing cognitive enhancers for the sake of some therapeutic purpose – but the point stands that this is a long way from having identified a case for cognitive enhancement across the board. It is certainly not an indication that we should be looking to make use of cognitive enhancers as a means of ensuring the unfettered growth of a technocentric world.

## The social benefits of cognitive enhancement

These lines of argument have concentrated on what one might call the 'first person' aspect of enhancement – the benefits that we might hope would accrue to the enhancee in a fairly direct way. But it's possible that we might want to endorse the use of cognitive enhancements for other reasons, and that these might contribute indirectly to the flourishing of the enhancees.

Writing on the *Practical Ethics* blog in 2010, for example, Julian Savulescu suggested that

> if one takes those people in the top 1% of the population of IQ, the top quarter of that top 1% produce more than twice as many patents as the bottom quarter. So even if you are in the top 1%, enhancing your IQ might enhance your creativity and inventiveness. It is important to recognize that cognitive enhancement is an important social and economic issue.[21]

This claim reasserts one made by Persson and Savulescu in 2008 – that cognitive enhancement would serve to accelerate this development:

> If, say, the mathematical gift of Euclid and subsequent mathematicians had been biomedically or genetically enhanced, the human species would presumably sooner have reached the state of contemporary mathematics and, consequently, hypothetical present day mathematicians would have reached a state of the art that actual present day mathematicians cannot divine.[22]

It is not impossible that enhanced mathematicians may be able to do things that are simply beyond the capacity of unenhanced humans; but

> the most likely effect of cognitive enhancement by genetic or biomedical means may be to speed up a growth in knowledge that would otherwise have taken humanity a longer time to achieve.[23]

And as for mathematics, so for other sciences. There are problems that face us today in, medicine, or meteorology, that we may not solve for generations at our present rate – if we are capable as a species of solving them at all. But were we to embrace enhancement, we may find that those problems that we can only solve slowly today become problems that we could solve quickly; and those problems that we cannot solve at all become problems that, sooner or later, we will.

The benefits appear clear. A humanity that understands the weather and climate better would be less vulnerable to at least some natural disasters. It might be possible to put a brake on anthropogenic climate change through climatological engineering; but even if that is impossible, it would be of great benefit to have a better idea of what will happen so that changes in crop patterns, public health measures, and so on, can be brought about the better to provide for the needs of humanity. By the same token, diseases whose cure currently presents us with a very hard problem might turn out to be much less tricky to deal with in a world where the top medical scientists are 10 per cent cleverer; cheap, clean, and renewable energy might be attainable (or attainable sooner, and more easily) if we were just a little smarter; and so on. Since we all have a reason to want diseases to be cured and cheap, clean, renewable energy to be supplied, we have a reason to support the means to achieve those ends – up to, and including, cognitive enhancement.

No less importantly, cognitively enhanced scientists might be better at spotting when they're on a wild goose chase, and so be better equipped to concentrate their energies on research that's likely to bear fruit. So it might be that we could present a case for cognitive enhancement being desirable on the grounds that it will make it more likely that beneficently-inclined people will be able to make the world a better place efficiently.

Before going further, it's also worth noting that Savulescu and Persson, in their 2008 paper, also worry that cognitive enhancement may have detrimental effects without ensuring moral enhancement – the promotion of pro-social behaviours – first.[24] The gist of the problem they outline is that scientific development has already made possible the prospect of weapons of mass destruction that could cause considerable harm; and expanding and accelerating this development will make it more likely that these weapons get into the hands of an increasing number of people. And

> if an increasing percentage of us acquires the power to destroy a large number of us, it is enough if very few of us are malevolent or vicious enough to use this power for all of us to run an unacceptable risk of death and disaster. To

> eliminate this risk, cognitive enhancement would have to be accompanied by a *moral* enhancement which extends to *all* of us since such moral enhancement could reduce malevolence.[25]

I shall return to the question of moral enhancement in Chapter 7. However, their worry does not seem to be completely off the wall: if we have the ability to make every Watson a Sherlock Holmes, we presumably also have the ability to add to the villainy of Moriarty the wit of Irene Adler, with the attendant undesirable consequences.[26] Yet these worries can be mitigated. After all, if we can improve cognitive abilities, there is no reason to suppose that we couldn't make every Sherlock a Mycroft; and, anyway, the possibility of undesirable consequences may count as *a* reason not to do something, but it need not count as a sufficient reason. Many things have potentially undesirable consequences, but enough desirable ones to make them worth doing on balance all the same.

Does an appeal to scientific advancement, then, generate a strong reason to pursue cognitive enhancement? Not obviously. The argument has a number of vulnerabilities.

Let's begin by examining the appeal to the increased capacity for invention and innovation that would supposedly derive from cognitive enhancement. It's only superficially strong. If the production of patentable work is our yardstick on this front, it may well turn out that such production has *something* to do with cognitive ability; but the relationship is nevertheless much more complicated than is allowed for in that account alone. For one thing, simply noting the number of patents generated is quite a crude measure, since many (and maybe most) patents do not recognize major break throughs, but minor improvements or variations of existing inventions – and many other more ambitious patents come to nothing. This suggests that the intellectual ability of the patent-holder doesn't have to be titanic, and that a super-brain may not be all that important: the mere luck of having the opportunity to file a patent detailing a minor modification to some item or technique would count for a lot. This kind of luck is not enhanceable.

At the same time, real progress often comes in forms that cannot be patented as insight is slowly accreted – and so, even if this accretion were to be made faster by cognitive enhancement, patent figures would not be the best measure of success. But it is not obvious that the right kind of accretion would be much helped by enhancement anyway.

Besides: even if it's true that the top 1-in-400 (that is, a quarter of the top 1 per cent) produce such a disproportionately high number of patents, that's plausibly because they're the top quarter of a percent, rather than because they're extra bright. Patents and honours gravitate towards some people not so much because of their intellectual capacity as measured in absolute terms, as because they are cleverer than the others – rather in the same way that the gold medals and records gravitate towards athletes who are faster than the others, irrespective of whether they are particularly fast in absolute terms.

Dismissing a suggested measure of scientific development isn't the same as showing that there won't be any. But there are reasons to be sceptical about whether cognitive enhancements would make all that much of a difference to the speed at which major breakthroughs are made, because it's not a given that high cognitive ability is of central importance. Notably, massive institutional support is likely to be a factor in a good portion of the kind of work that makes a difference to human flourishing, patentable or otherwise, and the same applies to an even greater extent to Nobel or Nobel-equivalent work: though it's still just about possible for hobbyists or non-professionals to function at this level in areas such as mathematics – Grigori Perelman provides an example of this – it's very unlikely indeed. The fact that people like Perelman do exist demonstrates that institutional support is not a *necessary* condition for great progress; but the fact that so few people like him exist is at least a sign that it's a significant factor. Moreover, the institutions that play such an important role in supporting the work that create significant breakthroughs are likely to be the ones that attract the best brains anyway – or at least to have the ability to sift out the less-than-best – so there'll be a natural agglomeration of talent there. And this counts against there being much of a difference made by cognitive enhancers for another reason.

The agglomeration of talent in one place is likely to create a dynamic that increases the likelihood of significant work being done by anyone at all; it is unclear that one would need enhancements of the bioscientific sort to have one's intellectual performance boosted by being a part of a close 'community of talent': being part of that community is an enhancement in its own right, and there is no particular reason to think that chemical or neurological enhancement would add anything that such membership couldn't.

Vitally importantly, we should not forget that being a member of a community of talent need not mean being one of a group of corporate employees; the internet makes it possible to crowd-source intellectual power. For example, the website Kaggle describes itself as 'a platform for data prediction competitions'; effectively a forum for members of the public to try their hand at solving problems, it

> allows organizations to post their data and have it scrutinized by the world's best data scientists. In exchange for a prize, winning competitors provide the algorithms that beat all other methods of solving a data crunching problem.[27]

The website claims never to have failed to outperform a pre-existing accuracy benchmark resoundingly, explaining this in terms of the benefits of having several problem-solving approaches working at once, and in terms of there being inbuilt competition.

> Real-time feedback is given on a live leaderboard, so when somebody makes a breakthrough, others revise their own algorithms to outdo the leader's performance. This leapfrogging continues until participants reach the full extent of what is possible.[28]

If Kaggle's claims are to be believed, they represent a challenge to the idea that the means to improve the world could be had faster by enhancing individuals – or, rather, that enhancing cognitive ability has to be a matter of enhancing the contents of any one person's skull. It would appear that we can do better science, with world-improving outcomes, by facilitating interaction.

One might still suppose that enhancement techniques like drugs or direct stimulation of the brain would allow for a bigger boost. And – for sure – cognitive enhancement *might* mean that a certain number of discoveries might be made sooner. Lots of enhanced people logged into Kaggle might mean that it's even better than it claims to be right now. But whether they would be made significantly sooner, and in significant numbers, remains a moot point; and whether there would be any problems solved in a world with such enhancements that would remain intractable in a world without is dubious. For this reason, the idea that we would all flourish more in a world in which cognitive enhancement were the norm is one that is far from definite.

## The benefits of distraction

One other reason for scepticism about the supposed benefits of cognitive enhancement may be worth noting, too: it might turn out to be counter-productive.

Let's return to the example of the student who wants to be able to boost her exam performance. She is worried about being easily distracted and wasting time doing things that are not conducive to study; offered a drug that will increase her concentration, she takes it. In her example, increasing concentration may be desirable. The same kind of thing could be said for a surgeon, who will have a perfectly understandable desire to focus his attention on the operation he is performing or about to perform – his patient's life and welfare may depend on putting the needle in the right place, or making a very precise cut. Finally, we might imagine a researcher who believes that we can, and maybe ought, to develop new drugs to combat disease for the sake of making the world a better place as quickly as possible – and that this imperative could be better satisfied were researchers a little mentally sharper and were able to concentrate more on the work they are doing. In all these cases, something to aid concentration would perhaps be welcomed, and treated as a cognitive enhancer.

I am happy enough to accept that, in the student's or surgeon's case, where there is a specific task in hand with clear things that need to be done, a lack of concentration might be problematic (though there is a set of questions about the use of examinations as a method of academic assessment that refuses to dissipate). But in the case of the researcher, this is not necessarily the case. This is because it's possible that the moments of staring out of the window, or checking the internet for the umpteenth time, are not always bad things.

Anecdotally, all kinds of breakthroughs have been made by people whose mind was *not* on the job in hand. For sure, much work remains to be done even if the outline of a basically good idea does come to people while doing other things. But, still, there is evidence to think that such events are not confined to anecdote and urban myth.

For example, a paper published in 2003 asked people who had achieved some level of creative success – as measured in terms of having patented, published or exhibited something – to perform a test that involved sifting out irrelevant information, and found that they were less likely to disregard such information and less likely to concentrate on the ostensible task in hand:

> The results of these studies and analyses indicate a substantial and significant relationship between a variety of indicators of creativity and reduced LI.[29]
>
> . . .
>
> These results lend support to the theory that there may be qualitative (e.g., failure to filter out irrelevant stimuli) as well as quantitative (e.g., high IQ) differences in the processes underlying creative versus normal cognition. . . . Thus, a deficit that is generally associated with pathology may well impart a creative advantage in the presence of other cognitive strengths such as high IQ.[30]

Put into layman's terms, what the study indicates is that creativity is associated with intelligence, but also with what the writers call a 'failure to filter out irrelevant stimuli' – that is, a tendency to get distracted. Such studies should be treated with care; but if there is anything to the suggestion that there's a correlation between creativity and having a wide 'mental sieve', increasing concentration – tightening the mesh of the sieve – may not be the best way to solve some problems. The researcher may find that extra concentration helps him perform routine tasks and follow established protocols, but that what is required to solve new problems is the ability to think creatively in a manner that is not easily accommodated by those protocols.

So it's possible that what counts as an enhancement in one context might count as a retardant in another; a person may benefit from boosted concentration when it comes to finessing an idea, but there might not have been the idea to begin with had she been concentrating too hard. And finding a way to reduce latent inhibition might be a good way to generate ideas, but might get in the way of translating those ideas into anything concrete.

Is this such a worry? Couldn't we fairly easily imagine a world in which people recognize that sometimes there will be a need to concentrate, and sometimes there will be something to be said in favour of free-association, and in which they adapt the enhancements they use accordingly – sometimes taking drugs to focus, and sometimes taking drugs to dream? Of course we can. But the problem is not that different drugs may be required at different times, but that people may not be very good at knowing which would be the right drug to take. For example, as I write this, I'm half-listening to the sounds coming from outside my window, thinking about the phone call I need to make later and

the train journey I need to make tomorrow, and wondering why a patch of my lawn has died. I'm also thinking about the couple of chapters that are currently proving more difficult than expected to get onto the page, and about the missed publisher's deadline. It's tempting to think that if I could mute this mental noise, I would write more efficiently and clearly. But even when I'm doing or thinking about something else, it's possible that problems I'm encountering in building the argument are simmering away somewhere on a mental back burner, and that that is the best way for them to be solved or avoided altogether. Whether this possibility corresponds with reality is hard to tell; but unless we can tell, there doesn't seem to be much point in taking a drug to enhance the characteristic we think we need to accentuate; we might have chosen wrongly.

## Alienation revisited

It's worth revisiting the 'Argument from Alienation' in this context of a consideration of the alleged social benefits of cognitive enhancement. It might not be desirable to promote cognitive enhancement as a means to generate welfare, because in doing so, one risks treating the would-be enhancees as functionaries of an overriding concern for welfare. It might be that people like Savulescu and Persson are correct to suspect that a world in which people are cognitively enhanced will be a world in which there are more good things, more quickly generated. For the sake of the argument, let's let that pass. However, there remains the possibility of embracing these good things as a desirable spin-off from enhancement, rather than as a reason to enhance in their own right. That is: there's a moral difference between a world in which people are cognitively enhanced for their own benefit, and – as a consequence that is foreseen and welcomed but not directly intended – happen to be in a position to generate a range of social benefits as well, and a world in which it's the social benefits that are in the moral driving-seat; and even if Smith himself is one of the people who benefits in the long run, this probably oughtn't to diminish the worry too much: what counts is that Smith would be a component part in the welfare-generator.

It may be harder for Smith to flourish in the face of the expectation that he will generate benefits. Indeed, in the face of this expectation, he may be less motivated to do anything much at all.

## The case for cognitive enhancement: Not wholly proven

As I noted at the start of the chapter, humans have been making use of cognitive enhancers for thousands of years; some of these enhancements are as simple as the proverbial knotted handkerchief, and others are as complicated as a

supercomputer. The reason these inventions have caught on is not, I would contend, because they help us make the world a better place – although they might. The reason they have caught on is the same as the reason why threshing machines have caught on among agriculturalists: they make life easier. From the point of view of securing and maintaining as good a life as possible, this is important. A person who has to labour hard for that outcome is, fairly straightforwardly, living a less good life than someone who has to labour less hard for the same outcome.

There is a role for cognitive enhancers in improving our lives, just inasmuch as *some* kind of cognitive work has to be done, and if something is available that will make that work easier, then so much the better. What is less clear, though, is that there is much of a role for 'neuromanipulation' as a means of enhancing the functions of the mind. We already know how to outsource a good many of the intellect's tasks, either to people or to things; the idea that bioscience may be a way to ensure flourishing is not obviously one that we have to accept. That is to say: a would-be cabbie may – let's say – find that his life is improved if some stimulus can be applied to the hippocampus, a region of the brain that has been found to be larger than average in London taxi drivers.[31] On the other hand, it would also be improved by buying a satellite navigation system. And while a satnav might mean we admire the knowhow of cabbies slightly less, it wouldn't remove anything all that important (and, anyway, the job might still require *noûs*). The same applies in other walks of life.[32]

More tentatively, we might wonder, too, about whether 'workaday' improvements in cognitive functioning would really be necessary to make life better – and whether their promotion might make life worse. I noted at the start of this chapter that lapses in memory – as when we forget a colleague's name – might be embarrassing, and make some kind of enhancement attractive. One might be tempted to wonder whether at least some of the badness comes from embarrassment; but if Farah et al. are accurate in saying that this sort of lapse becomes more common from the fourth decade of life onwards and is completely unremarkable, one way of mitigating that badness is to remove the embarrassment by simply admitting to ourselves and others that we aren't as quick as we once were. Promoting enhancers by drawing attention to these small errors could in fact contribute to the creation of a problem that the enhancers are supposed to solve. Lapses in memory may threaten a dip in flourishing; but being ashamed about them, when they are so trivial, may threaten something equivalent.

In some circumstances, some form of artificial improvement of intellectual capacities by bioscientific means may be desirable. Yet what seems clear is that the desirability of boosting people's brain power depends on us working in a certain precise context – surgery, or revision, or something of the like. This isn't enough to back up the claim that cognitive enhancement is generally all that desirable or worth the effort. Cognitive enhancement may well have benefits to

those who struggle with the world, but it doesn't follow from that that it will benefit everyone; and if the benefits are more widely distributed, it still doesn't follow that bioscientific means would be the best way to realize them. It is quite possible that they wouldn't be. I advance as a hypothesis that, by being satisfied to farm out our cognitive processes to books, computers, and other people, we would lose no significant benefit.

## Notes

1 Bostrom and Sandberg, 2009, p. 311.
2 Ibid., p. 312.
3 Ibid., p. 321.
4 Mithoefer et al. 2013, p. 36; cf. Sessa, 2007, *passim*.
5 see Plato, 1952, 274c ff; 1981, *passim* (but esp. 81b-e and 85d-e); but cf. Plato 1999m 142d ff. vide also Havelock, 1982, pp. 41–7.
6 Farah et al. 2004, p. 422.
7 Wittgenstein, 2001, § 265.
8 Wagenaar, 2008, *passim*.
9 Vide Vedder and Klaming, 2010, pp. 29–30.
10 'Funes, the Memorious', in Borges, 1998, p. 104. For Plato, an understanding of the idea of the dog-in-general would require remembering the Forms; for Borges, it would require forgetting.
11 Mayer-Schönberger, 2009, p. 21.
12 Parker et al. 2006, p. 46.
13 Ibid., p. 35.
14 Ibid., pp. 38, 42; slightly modified.
15 Mill, 1998, p. 140.
16 Ibid.
17 Bostrom and Sandberg, 2009, p. 324.
18 Ibid.
19 see Marx, 1970, *passim*.
20 see 'The Question Concerning Technology', in Heidegger, 1999, passim.
21 Savulescu, 2010, slightly modified.
22 Persson and Savulescu, 2008, p. 165.
23 Ibid., p. 166.
24 Ibid., *passim*; cf. Persson and Savulescu, 2012, *passim*.
25 Persson and Savulescu, 2008, p. 166.
26 Irene Norton, *née* Adler, is a character in the story 'A Scandal in Bohemia', and one of the few people to best Sherlock Holmes. Holmes refers to her 'always under the honourable title of *the* woman'.
27 http://www.kaggle.com/.
28 http://www.kaggle.com/about.
29 LI stands for 'latent inhibition', which is the tendency to pay more attention to novel stimuli than familiar ones.

**30** Carson et al. 2003, p. 505.

**31** Parts of the hippocampus have been found to be enlarged in London taxi drivers who have done 'the knowledge' (see Maguire et al. 2000); this kind of phenomenon would seem to suggest, first, that the brain is fairly plastic, and second, that experiences can shape it. However, other work would seem to indicate that these gains come at a cost: taxi drivers appear to have a lower ability to acquire *new* visuo-spatial information than do bus drivers (Maguire et al. 2006).

**32** cf. Bostrom, 2008, p. 30.

# 7

# Good Is as Good Does? The Case of 'Moral Enhancement'

## The possibility of 'moral enhancement'

Over the past couple of chapters, I've tried to outline a position that is sceptical of the ability of bioscientific advances to make all-things-considered good enough lives enormously better – but none of this has amounted to hostility. However, in respect of one potential and putative enhancement, I think that there is room to add to a general scepticism a more positive reason to reject methods to alter human capacities: this is in relation to so-called 'moral enhancement'.

Very broadly, the term 'moral enhancement' capitalizes on our increased understanding of the neurobiological factors that determine behaviour, and puts forward the idea that there could be interventions – most likely, by altering the brain or some other determinant of behaviour – that would make a person more moral: that is, more disposed to do the right thing. In this vein, Tom Douglas' formula for moral enhancement is that

> [a] person morally enhances herself if she alters herself in any way that may reasonably be expected to result in her having morally better future motives, taken in sum, than she would otherwise have had.[1]

This account may well invite criticism – but it'll do for now.

Of course, there is a great deal to be said for moral enhancement taken in the broadest sense: parents teaching good behaviour to their children are engaged in what they at least hope will be a form of moral enhancement – and it's expected that they ought to be. And one aspect of being a moral agent at all is the expectation that we should seek to restrain our baser instincts and make ever finer distinctions when they are there to be made; this, too, is a kind of moral enhancement. For example: imagine that Bob has been brought up to dislike and distrust members of a particular ethnic group, and absorbs these attitudes with his mother's milk. Bob leaves home and goes to university, where he finds himself forced to associate with someone from that ethnic group, and it begins to dawn on Bob that this person really isn't as bad as all that; by the end of his first term, Bob has been forced to re-evaluate his attitudes. But it's hard to shake years of conditioning; and so while he is now more willing to accept the baselessness of his

old attitudes rationally, there is still a slight residue – an involuntary half-second reluctance to trust members of that group, say. However, Bob knows that this residual feeling is baseless and unjust, and tries to suppress it – to teach himself to react differently. This is, I'd contend, a kind of moral self-enhancement; and it would be not only admirable, but quite possibly mandatory.

So it would be absurd to present a case that moral enhancement is necessarily wrong. But there are still problems, as we shall see over the coming pages; and they might well be quite big.

## Strategies for moral enhancement

A paper by Tom Douglas in the *Journal of Applied Philosophy* from 2008, which is fast becoming the *locus classicus* in respect of moral enhancement, considers the case for intervening to ensure that people have better motives than they otherwise might have had. In Douglas' scenario, the intervention is by means of a pill, but we could imagine other methods such as direct brain stimulation of one sort or another; the point is that a person who previously had a propensity to behave in *this* way could undergo some kind of procedure that would increase the likelihood of his behaving in some other – implicitly more desirable – way. More, Douglas claims that there are some 'counter-moral emotions'[2] that would always be (though maybe 'provide' would be a better word) uncontroversially bad motives for action, and the attenuation of which would reduce their ability to interfere with good motives.

One example that he gives of a counter-moral emotion – both bad in itself, and likely to interfere with good motives that a person may have – is a strong aversion to certain racial groups. Another is a strong impulse towards aggression, which may occasionally be in order, but is likely in most cases to be a morally bad motive. But these examples are meant to be illustrative more than definitive; what counts is that

> there are some emotions such that a reduction in the degree to which an agent experiences those emotions would, under some circumstances, constitute a moral enhancement.[3]

Finally, he hypothesizes, inducing a reduction in the degree to which a person experiences those emotions is not impossible in principle. So, if we can control the expression of (say) a gene that is strongly linked to aggression, or make some difference to an agent's neurochemical balance such as to make him more likely to behave in a desirable way – it's widely believed that oxytocin, for example, can modulate behaviour to make it much more pro-social – then we have at our disposal a method of making people better. The details need not bother us all that much: what counts is the basic idea of intervening directly to make a person 'better behaved'.

One immediate worry here might be that we're embarking on some kind of *Clockwork Orange* scenario in which a rich moral life is displaced by conditioned good behaviour, with the implication that 'enhanced' goodness is no such thing. I'll return to a version of this worry later; but the objection may be a little simplistic in these terms. Importantly, Douglas talks a lot in the paper about a person enhancing *himself* – and, indeed, one possible use for moral enhancement of the sort that appears to be envisaged here is to help people like Bob in the example above control the emotional responses to certain situations that they would prefer not to have. It would still, in the end, be up to him to try to cultivate the right attitudes – it's just that the task would be easier; and, to reprise the Real Tomato Argument from Chapter 5, taking a pill need not mean that the results were void of moral merit.

Douglas' emphasis is on interventions on currently existing people; but we could also intervene to ensure that people have more desirable moral dispositions by making more positive decisions about who gets to be born in the first place. In a 2006 paper, Julian Savulescu et al. consider the relative merits of eugenic selection for desirable behaviour.[4] Like Douglas', their argument relies on the premise that at least a significant factor in the determination of moral behaviour can be naturalized – that is, that it can be attributed to an empirically scrutable fact about a person's physical make-up. They begin their argument by pointing out that there is a credible genetic factor in at least some morally problematic behaviour. For example, there is evidence that a genetic abnormality that leads to low levels of the enzyme monoamine oxidase A (MAOA) can be correlated to anti-social behaviour – and that this suggests that it ought to be possible to do something to reduce the levels of such behaviour. They conclude that there is a sound moral reason to select against embryos that carry the low MAOA mutation.

Though their preferred strategy is to select embryos rather than attempting to alter particular genes or the expression of particular neurotransmitters within a given embryo or person – they prefer genetic selection to genetic enhancement in this case – I think that it is still fair for the sake of this chapter's concerns to place embryo selection of this sort at least broadly within the scope of arguments for moral enhancement. They are still talking about interventions that draw on bioscientific insights and are expressly intended to improve behaviour; this is close enough to moral enhancement to make no real difference for my purposes.

Such enhancement – of the sort envisaged by both Douglas and Savulescu et al. – could be presented as serving entirely the interests of the society in which the embryo would come to birth, since it is reasonably safe to assume that society has an interest in at least most of its members being less aggressive. This impression might be furthered by the way that Ingmar Persson and Julian Savulescu talk about the need for moral enhancement in their 2012 book *Unfit for the Future*. They argue that the big problems that face humanity in the

decades and centuries to come and threaten the very existence of the species, or (at best) the welfare of its members – notably, climate change, but also other threats that arise from the fact that it is very much easier to kill great numbers of people by action or inaction than it is to save fewer – are problems because we have a systematic bias in favour of those who are spatially and temporally close to us, and because we have a bias to treat moral responsibility narrowly, in terms of causation.[5] They are problems that – it is hypothesized – will not be solved by technology alone, not least because there is currently little incentive to develop it. Understanding the neurological and genetic reasons for our behaviour, though, would mean that we would be in a much better position to do something about it, and to amplify and solidify 'the internalisation of moral doctrines' by tailoring our genes and neurology towards the ends required by morality properly and widely understood:

> We call moral enhancement by such means *moral bioenhancement*; possible examples of moral bioenhancement would be drug treatment and genetic engineering.[6]

What is meant by moral enhancement here is not reducible to making people more pro-social (although increased sympathy would come in to it), but about making them better at being good. Oxytocin is one suggested method; but the precise method isn't one that is all that important.

However, such interventions may turn out to serve the interests of the enhanced person and person-to-be just as much as the people around whom they live:

> Given that compulsive violent or criminal behaviours are typically punished with incarceration as well as social rejection, we can assume that these behaviours typically harm a person who expresses them, as well as their victims. Selections which constrain a person from making such harmful choices may therefore be justified.[7]

As such, there is a case to be made for it being the sort of thing that may make the enhancee's life better, at least indirectly. And so the case for bioscientific moral enhancement seems to be fairly straightforward: making people morally better is desirable just insofar as it is morally desirable to be morally good and to promote the good; it is desirable because those of us who have inclinations that we would rather not have might thereby have a strategy to cultivate our own virtue more easily; it is desirable because, if it is possible to reduce the frequency and likelihood of morally undesirable behaviour, that will probably mean fewer criminal and harmful actions, which means that there is a good chance that less harm will accrue to fewer victims; and it is desirable because fewer people will have to suffer incarceration, social resentment, and the other harms that follow from the misfortune of having a disposition towards aggression.

## The argument from freedom

There is a range of problems that might be raised for proponents of moral enhancement – and possibly with moral enhancement itself. John Harris has argued against moral enhancement on a couple of grounds. One of these is that avoiding the potential problems of human behavioural blundering requires not so much *moral* as *cognitive* enhancement – or, perhaps, that cognitive enhancement will allow us to avoid moral disaster much more effectively than 'moral enhancement' as outlined above would. I'll return to this theme in a little while. Another objection, though, is that moral enhancement programmes are antagonistic to other desirable characteristics – notably freedom.

It's a concern with freedom that brings Harris to see moral enhancement as a source of positive suspicion. For him, citing Milton, 'sufficiency to stand is worthless, literally morally bankrupt, without freedom to fall'[8]; he claims that

> ethical expertise is not "being better at being good", rather it is being better at knowing the good and understanding what is likely to conduce to the good. The space between knowing the good and doing the good is a region entirely inhabited by freedom. Knowledge of the good is sufficiency to have stood, but freedom to fall is all. Without the freedom to fall, good cannot be a choice; and freedom disappears along with virtue. There is no virtue in doing what you must.[9]

and that

> the sorts of traits that seem to lead us to wickedness or immorality are also the very same one required not only for virtue but for any sort of moral life at all.[10]

Harris does not seem to be denying that the targets that people like Douglas or Savulescu or Persson have in mind are anything other than worth pursuing – or even urgent. But what is at stake is the specifically *moral* aspect of achieving them. Grant that the manipulations envisaged will make people behave in the manner that morality would dictate: there would be no actual *morality* on display, since agency itself would be diminished; we would have moved towards the mere playing-out of a pre-determined set of actions that would not indicate any particular praiseworthiness on the part of the people carrying them out. Those people's 'moral' capacities would not be enhanced at all – we would not see any improvement in morality richly understood. (The distinction is rather similar to that drawn by Kant between acting *from* duty, and acting merely *in accordance with* it.[11])

Are these worries sufficient to undermine arguments for moral enhancement? In the 2006 paper mentioned above, Savulescu et al. appear to be prepared at least to acknowledge their coherence: they admit that appeals to the importance of an 'open future' may generate an argument against their behavioural genetic selection – although they also (rightly, in my opinion) note that open future arguments are not powerful.[12] But in *Unfit for the Future*, Persson and Savulescu

take issue with the idea that being biologically determined to have certain dispositions makes any difference to freedom. For example, suppose it is true that women are, by nature, likely to be more altruistic and less aggressive than men because of something to do with biology. (Cordelia Fine points out that the evidence backing up this claim is, at best, shaky[13] – but we can let that pass for the sake of the argument.) Well, even so, this does not mean that women are less free than men. If moral behaviour is straightforwardly neurophysically[14] determined, then all moral enhancement would do is tip the balance in favour of us doing praiseworthy things. We would be no more or less free than otherwise. Conversely, if moral behaviour is not neurophysically determined,

> [t]hen moral bioenhancement cannot be fully effective because its effectiveness is limited by the indeterministic freedom that we possess. So, irrespective of whether causal determinism or indeterminism reigns in the field of human action, moral bioenhancement will not curtail human freedom and responsibility. Biomedical manipulation cannot chance the basic laws of our behaviour by making us more (or less) causally determined . . .[15]

While this might answer some of Harris' qualms, his view of freedom still seems to be one in which freedom is magnified the more options that one has; hence to have one's brain manipulated in such a way as to make it much less likely that (say) one will become violent is to reduce freedom. But this prompts an important question: Is it true that the scope of someone's freedom really is related – at least directly – to the number of options open to her? If it is not, removing undesirable options would not necessarily be a constraint on freedom or moral action richly understood; and I'll argue in the coming paragraphs that reducing the number of options is *not* a constraint on morality or moral action richly understood – and that the 'freedom to fall' is therefore inessential to moral freedom *per se*. In fact, I think that having fewer options can be, morally, a good thing: and in a sense, it's what moral education is all about.

## Freedom and options

Suppose you are walking along a quiet street one night, and notice that ahead of you is someone rather substantially the worse for wear from drink. Strictly speaking, one of the things that you could do in this situation would be to attack and rob that person. There is nothing really stopping you behaving in this way. However, things can be *formal* possibilities without being *live* possibilities. A reasonably decent person would not see the option of attacking the stranger as a live possibility; it's simply not the kind of thing that appears in the suite of plausible things that could happen. This is not just because of worries about getting caught and preferring not to be punished: pertinent as those worries might be from one perspective, the truly morally developed person doesn't

even have to make that kind of calculation. Attacking people on the street is, in everyday parlance, unthinkable. To borrow a phrase from Bernard Williams, someone who has to think about whether attacking strangers on the street is permissible, weighing up the reasons to act against those not to, has had – and even if he decides not to act – 'one thought too many'[16]; and this is what I mean by the idea that there are certain formal possibilities for action that are not, in a morally well-constituted subject, real or live possibilities.[17]

Though writing in a different context, John McDowell makes the point that in moral upbringing, 'one learns not to behave in conformity with rules of conduct, but to see situations in a special light, as constituting reasons for acting'[18] – or, as in the example above, not acting. Thus, in an important passage, he claims that

> the attractions of whatever wickedness might bring do not constitute some reason for wickedness, which is, however, overridden by the reasons against it; rather, given that they are achieved by wickedness, those attractive outcomes do not count as reasons at all.[19]

Of course, street attacks – and worse – do happen; but while we condemn them and the sort of person who carries them out, I think it reasonable to say that it would not be unjust to think that a person who has to work out that and why she should not attack people on the street is morally lacking in some important way too. This is as much as to say as that a person with more live options for action is not always more morally praiseworthy than someone with fewer; having fewer options might, in fact, be a sign of being a complete moral being. David DeGrazia echoes this thought – satisfyingly enough, as part of a defence of moral bioenhancement – when he points out that '[a]bility to act otherwise, or freedom to fall, is not a necessary condition of free action'.[20]

Neither is it obvious how the inability to see certain logically possible courses of action as live options makes us less free, any more than the inability to see red as green does. But even if, for the sake of the argument, we *are* less free in some sense, it does not obviously follow from that either that a diminution of freedom for the sake of better behaviour is undesirable (DeGrazia again: 'Presumably, freedom has significant value. It doesn't follow that its value overrides everything else that matters'[21]) or that our action has any less moral merit or that we are in any way less whole moral agents. After all: not attacking people in the street does not really attract a great deal of praise in any case – that's simply a matter of behaving in a way that's minimally acceptable, and it's strange to think that our integrity as agents depends on there being a realistic chance of us doing what we ought not do and would not think of doing. I have any moral merit at all, it's more bankrolled than bankrupted by my utter disinclination to go raping and pillaging my way around Manchester.

We might even want to take things a little further, beyond simply denying that having fewer live options is compatible with moral agency properly understood: it is not completely *outré* to wonder whether having fewer live options could itself count as morally enhancing, such that the less quick to anger, the less aggressive, and so on that we are – the less we are slaves to our passions – the freer we are. At least some elements of Kantianism are sensitive to this notion of being free from the passions – but it's a line of thought that goes back as far as Plato. Having certain base impulses purged by moral enhancement of one sort or another may make us *more*, not *less*, free. So the idea that morality properly understood is corroded by attempts to make it less likely that we will choose to behave badly seems to me not to have much merit.

(For the record, it's worth noting the oddity that Harris should be keen on preserving the 'freedom to fall' in this context, but seems less devoted to it in others. In relation of children participating in – or perhaps 'being participated in' would be a better turn of phrase – scientific research, he and Søren Holm suggest that

> [o]ne of my main interests is my interest in taking myself seriously as a reflective moral agent, and my interest in being taken seriously by others. Identifying my moral obligations, and acting on them, is not contrary to my interests, but is an integral part of what makes me a moral agent. But what about children?
>
> If children are moral agents, and most of them, except very young infants are, then they have both obligations and rights. It is, furthermore not a requirement for having a moral obligation that a person is able to give a full account of the moral justification for that obligation.
>
> If a parent does not take this into account when making decisions for the child that parent displays one of the following attitudes; either the attitude that the child is not (and need not be) a serious moral agent at all, or the even more problematic attitude that the child is so deeply fallen in moral turpitude that it is not willing to discharge any of its moral obligations. The only way of taking the child seriously as a developing moral agent, interacting with other moral agents in a complex world, is to assume that the child is serious (or would wish to be serious) about discharging its moral obligations. Parents are therefore justified in assuming that the person they are making decisions for is a moral person, who wants to discharge his or her moral obligations.[22]

(This appears to boil down to the idea that we take children's moral agency seriously by making certain decisions on their behalf – which seems to be a claim that we can and maybe should deny their freedom to fall *as a part of taking their agency and capacity for virtue seriously*. I do not find this argument at all convincing – and Barry Lyons' dismissal of the idea is quick and cutting, and explodes it in a matter of a few sentences[23] – but, still: it's tempting to wonder just how much freedom to fall is necessary even by Harrissian lights.)

Having said this, I think that Harris is arguing on the right lines; his position can be fortified, and that's a task to which I'll turn my attention now. The first step is to invoke the spirit of Aristotle.

## Nicomachean moral enhancement

Aristotle is suspicious of the idea that one's moral characteristics are what they are by nature; he states at the start of Book II of the *Nicomachean Ethics* that moral goodness 'is the result of habit', which means that

> none of the moral virtues is engendered in us by nature, since nothing that is what it is by nature can be made to behave differently by habituation.[24]

Since a person's characteristic patterns of behaviour often can be changed, that means that the disposition to display certain or all virtues must be plastic. For Aristotle,

> we become just by performing just acts, temperate by performing temperate ones, brave by performing brave ones . . . It is the way that we behave in our dealings with other people that makes us just or unjust, and the way that we behave in the face of danger, accustoming ourselves to be timid or confident, that makes us brave or cowardly . . . Hence we must give our activities a certain quality, because it is their characteristics that determine the resulting dispositions. So it is a matter of no little importance what sort of habits we form from the earliest age – it makes a vast difference, or rather all the difference in the world.[25]

Habituation means that we give ourselves a 'second nature' that will determine how we act in everyday life; it is by emulating the best that one improves oneself as their habits and dispositions will become one's own. And in some cases, one doesn't even have to decide to emulate the best, the decision having been made by others:

> Legislators make their citizens good by habituation; this is the intention of every legislator, and those who do not carry it out fail of their object.[26]

The point is that Aristotle seems to have some account of a programme of becoming more virtuous that one could adopt or even impose on others. At least on the face of things, this looks as though it could be read as being a kind of claim about moral enhancement. Moreover, his account does not jar too badly with a number of modern intuitions about ethical formation; I'll come back to them in a little while.

In the Aristotelian picture, there is a meaningful parallel between acquiring the virtues – a disposition towards temperance, justice, and so on – and acquiring any other skills. Thus

> the causes or means that bring about any form of excellence are the same as those that destroy it . . . for it is as a result of playing the harp that people become good or bad harpists. The same principle applies to builders and all other craftsmen. Men will become builders as a result of building well, and bad ones as a result of building badly. Otherwise there would be no need of anyone to teach [people]: they would all be *born* either good or bad. Now this holds good also of the virtues.[27]

There is a slight oddity about the idea that one becomes a good builder by building well – we might ordinarily suppose that it's only good builders who are capable of building particularly well to begin with. But this is not a serious problem for my purposes – we could avoid it by rephrasing things slightly while retaining the gist – and his statement has a certain attraction to it. The point is that one would acquire practical skills by associating with skilled people and doing as they do, gradually becoming habituated to that particular skilled activity; likewise, one would acquire the virtues by associating with virtuous people and doing as they do, gradually becoming habituated to virtue oneself.

So let's accept that there is a mechanism by which one becomes virtuous or skilful, and take the example of a novice builder. He will be less-than-excellent to begin with; but we might reasonably hope that given a decent mentor and time, he will become excellent. If the development of his skills is retarded, he will remain a bad builder. The alternative makes little sense: being a bad builder is not a vocation in its own right; indeed, though one sets out to be a good builder, being a bad builder for a while is probably an integral part of realizing this ambition. Still, this does not mean that a person has any latent skill as a builder that habituation realizes. Before starting to build for the first time, a person is – to paraphrase Wolfgang Pauli – *not even bad*: not a builder at all.

It's tempting to give a similar account in respect of moral development. We tend not to blame children for bad actions, and this is because we think that, at least in some respects, they are not moral agents to begin with: they, too, are not even bad. The first stage of moral training involves the task of making a person – usually a child – capable of bearing moral epithets, just as the first stage of making someone a builder is to make him capable of bearing the appropriate epithets. At the same time, good behavioural patterns are instilled by socialization: getting children to write thank-you letters, for example, helps generate desirable patterns of behaviour that mean that courtesy comes naturally in later life, just as repeatedly building walls means that building them well soon enough becomes a matter of instinct.

The flip side of this is that Aristotle also says that there is a process by which one becomes a bad builder, harpist or whatever; and that it is the same as that by which one becomes a good one. If acquiring practical skills is truly analogous to acquiring the virtues, this would suggest that the manner in which someone becomes vicious is, in the first place, by habituation, and, in the second, like that in which someone becomes a bad builder.

The first of these ideas – that one can become vicious by habituation – seems reasonable enough, and seems to fit in respect of people who turn out to be moral failures; there're plenty of stories that we could tell in which a person seems to have become vicious at least partly because of some kind of malhabituation, and we do tend at least to be sympathetic to the idea that people can be habituated towards evil by a bad upbringing – which seems

to be the correlate of a belief that a good upbringing is a requirement for the formation of an admirable moral character. For example, the news in late 2009 and early 2010 carried the story of two young brothers, aged 10 and 11, tried for a brutal attack on two other boys of about the same age. The boys were found guilty, and such was the court's assessment of the danger they posed to others that they were detained for an indefinite period. However, the court heard how the defendants had been raised in what social workers called a 'toxic' environment at home – they had witnessed domestic violence, pornography, horror films, had access to alcohol, and had been fed cannabis by their mother.[28] Though this one rather lurid case is clearly not enough to sustain a whole argument, it is illustrative of the idea that behaviour is in at least some respect associated with upbringing – that a tendency towards violence is not too surprising in someone who is habituated to violence from an early age. This looks to be a good example of the process that would habituate or engender moral virtue being the same as that which would habituate or engender moral vice. The process that can create virtuous dispositions can also destroy them.

So far so good: the problems begin when we ask more deeply about the analogy between becoming vicious and becoming – say – a bad builder. The habituation thesis is definitely not a plausible account of the process involved in becoming a bad. If it's true that becoming a bad builder is a step on the way to becoming a good one, then it won't destroy the excellences associated with building well, or make it more difficult to become good – which is quite unlike moral development, where acquiring a vice is antagonistic to the corresponding virtue. Similarly, while it might be possible to desensitize a previously virtuous person to vice, or *vice versa*, it is not obvious that you can desensitize people who possess other practical excellences. For example, a virtuoso harpist may spend a great deal of time in the company of barely competent students, but while exposure to them may mean that she understands their mistakes, that will not destroy her excellence. To lose that excellence would require that she fall out of practice altogether. By contrast, whereas a virtuous person who is rarely presented with any particular requirement to exercise his virtue will not obviously be less good at doing so when the need arises, he may have his moral sentiments coarsened by spending a great deal of time with brigands. This means that the process of becoming a bad person is rather different from the process by which one becomes a bad musician, builder, or whatever.

It is possible to salvage the analogy, provided that we're prepared to accept that there are two distinct sides to Aristotle's idea, which Aristotle has mistakenly run together. One of these sets is 'functional', and relates to the ability to do well; the other is 'cognitive', and relates to the ability to judge well. (The ability to judge well may be required to do well, but it is not the same.) This distinction will turn out to have a role in thinking about biomedical moral enhancement, and in rebuilding something like an argument from freedom.

## Rebuilding the argument from freedom

Let's return to the builder example. A good builder will be able to plaster a wall evenly, and will take proper pride in a job done well. He will also be able to judge what kind of plaster to use, how much water is required, and so on, and none of this is required actually to put the stuff on the wall. Similarly, a convincingly good agent will not only have a desire to do the just thing and a horror of doing the unjust; he will also have the entirely separate cognitive skills required to assess what justice demands.

Now: in his account of moral enhancement, Douglas is concerned explicitly about instilling better motives in people, and not so much about processes of moral reasoning. The concern appears to be attributable to the way that understanding of the phrase 'moral motive' has to do with the 'psychological – mental or neural – states or processes that will, given the absence of opposing motives, cause a person to act'.[29] To this extent, his paper speaks to the functional aspect of Aristotelian ethical formation – the bit that's analogous to spreading the plaster on the wall – more than the ability to judge what the right thing to do would be.

But we need to be careful with the language here. To be caused to act in a certain way is not the same as acting for a cause. Douglas is talking about the things that cause a given behaviour; the cause in this sense is whatever it is that explains and brings that behaviour about. But this is not the same as talking about the cause that a person may have; to use the word 'cause' in this sense is to talk about whatever it is that invites behaviour. A drone aircraft flying above enemy territory can be caused to drop a bomb; but it does not drop a bomb for the sake of any cause.

Put another way, he is talking about the reasons *why* an agent acts as he does, but not reasons *to* act. And the important thing is that the two things may come apart: a person's mental or neural states explain *why* he likes ice-cream, but do not furnish a reason *to* prefer it; facts about the evolutionary advantages of ascribing agency to observed phenomena may tell us *why* humans believe in deities, but they will not provide a reason *to* believe. In the same sort of way, we can coherently prise apart the concept of a motive from the concept of a motivation. To have a motive is to have a reason to act or a cause in the name of which one might act. To have a motivation, by contrast – to be motivated – is rather to have been the subject of a certain kind of causation. An agent who acts for a cause, or who has a reason or a motive to act, is an agent who has identified an impetus for acting that he has adopted as his own. By contrast, an agent who is caused, or whose actions are for a reason or are motivated by someone, may actually be in the dark about quite why he acted as he did – and may even be less able to answer the why question than an external observer would be. (We can imagine someone in the dock or the therapist's chair who only a long time after the event comes to understand why he acted as he did.

'Cause' in his sense would be an explanation for action, not an invitation to it.)

To this extent, I think that Douglas has perhaps fallen into a linguistic trap. When he talks about moral enhancement in terms of an agent's having

> altered herself in a way that may reasonably be expected to result in her having morally better future motives, taken in sum, than she would otherwise have had[30]

this is something that can be perfectly well be accommodated by something along the lines of the Aristotelian picture of moral formation by habituation that I outlined above. But when he suggests the possibility that this may be achieved by biomedical means, he has stopped talking about moral motives – he is now talking about moral motivations.

How does this help revivify the argument from freedom? Well, for Douglas, it would appear to be sufficient that an actor is caused by some set of mental or neural characteristics to behave in desirable ways. So imagine a variation on the Frankenstein story, in which Victor Frankenstein has the opportunity to instil in his monster certain dispositions and inclinations: to make him, wherever possible, more responsive to desirable moral motives and less likely to behave in morally undesirable ways.

But this would fall foul of the argument from freedom. This is not so much because the monster would have fewer options, as because it is caused to act without acting *for* a cause: freedom attaches to adopting a cause, having a reason, or having a motive. It's for this reason that McDowell's virtuous agent is no less free for being unable to see the other (merely) logically possible courses of action that his *Doppelgänger* on Twin Earth may have chosen. Though there is a sense in which Douglas' agent and McDowell's are alike, inasmuch as that acting in a different manner is unthinkable, there is another sense in which the difference is vast, just because the McDowellian agent has a reason that makes the other reasons to act inert. Moreover, he could, if things turn out in the right way, morally re-habituate himself to some other cause – a point that will be important in a little while. This means that he would be, in the Miltonic sense, sufficient to stand, though free to fall.

It might even be possible to rehabilitate the concept of dignity in this context. Notwithstanding the familiar problems with the concept, if anything stands a chance of eroding an agent's dignity, this does: in effect, we would be trying to make them beholden to our moral preferences. And even if this is what parents do as part of the normal course of raising children anyway, those children are able to adopt the moral shape of their upbringing as a cause of their own.

Already, then, we've got a reason to wonder whether moral enhancement is compatible with the good life. If part of being a fully-fledged human moral agent is to be morally free, it could turn out that moral enhancees are missing out on a significant dimension of the complete human life. And to be caused to

act in accordance with someone else's preferences is no less corrosive of a good life, richly understood.

Does the same apply to attempts that we might make to enhance ourselves? Imagine that someone is told by others that his habitual patterns of behaviour are morally unacceptable, and resolves to alter that behaviour. There does seem to be less of a worry about dignity here, for although our anti-hero is attempting to force himself to meet others' expectations, that attempt is *his*. So a person who tries to 'recondition' himself to be more likely to behave in accordance with what his better angels would prefer, and takes full advantage of the insights provided by brain-science to ensure that he succeeds in altering his behaviour, might have more of a reason to do so – and would not necessarily have to worry about eroding his own freedom or dignity. (Having said this, Harris is not wholly sold on the idea that this would be so straightforward – moral bioenhancement, he worries, 'may well dull [moral consciousness] to the point where the individual is no longer choosing'.[31] On the basis of what I've been arguing so far, the absence of choice – conscious choice at least – doesn't worry me too much; but the reason why conscious choice is absent does matter; and inasmuch as that Harris is nodding towards the distinction between being caused and having a cause, I am in agreement with him. And self-administered moral enhancement that elevates the importance of being caused in comparison to that of having a cause might quite straightforwardly be incompatible with human flourishing, since it could quite straightforwardly be said to disrupt an aspect of complete human functioning.)

Still, there would be other problems to be overcome all the same. Freedom isn't the only concern about moral enhancement. The big problem in respect of attempting to give someone better motivations is that there is room for reasonable disagreement concerning what 'better' would be.

## The argument from reasonable disagreement

In a good deal of the literature, moral enhancement is taken to imply making people more pro-social in some sense – I count Persson and Savulescu's reinforced altruism across space and time to be a version of increased pro-sociality. But since 'morality' is logically different from 'pro-sociality', there is obviously a need to provide an argument about why moral enhancement should be taken to mean boosting pro-sociality. And it makes sense to wonder whether, in the absence of such an argument, we ought to be wary of 'improving' an agent's dispositions until we are sure of the measure by which we are supposed to assess that improvement. Call this the 'Argument from Reasonable Disagreement', or 'ARD' for short.

We might want to assume that there are certain traits the desirability of which should be minimally controversial (although even here, the existence

of a consensus doesn't foreclose the possibility of reasonable disagreement[32]). Douglas' own example of non-racism seems to fit into this category – although we should be wary of the possibility that, if reducing racial aversion this strikes us as being obviously worth doing, this is perhaps a warning that his example is too slight to carry the weight of a fully-fledged argument. But there are other dispositions whose undesirability is not so straightforward, and in such cases, there is a question to ask about what the next move should be. So, when Douglas also suggests that a propensity to violent aggression might also merit 'enhancing' away – and that's a trait that Savulescu et al. would happily eliminate as well[33] – the way does seem open for people like Harris and Chan to point out that we ought to slow down a bit, because there are some situations in which violent aggression might be either called for or, at least, not be wholly out of line. (Douglas admits as much himself.) Harris' preferred example here is of using violence to overpower a would-be hijacker, or someone else intent on causing damage or destruction[34]: violent aggression, channelled in the right way, might be precisely what's in order. While removing the disposition to violence may reduce the chance that someone will be a hijacker, it may also reduce the chance that people nearby would be able to do much to thwart a person sufficiently unenhanced still to mount a hijack.

Douglas' way of avoiding the pitfalls raised by this point is to rephrase the claim:

> [T]here are some emotions such that a reduction in the degree to which an agent experiences those emotions would, under some circumstances, constitute a moral enhancement.[35]

It's not clear that this is enough. For one thing, the claim now looks horribly formalistic: it doesn't tell us anything much about which emotions and dispositions we would do well to enhance or reduce. Neither is it clear that there would be agreement about at least some of the enhancements up for consideration. For there is, as yet, still controversy about the standard against which we might hope to calibrate our moral assessments.

For natural scientists, though there can be reasonable disagreement and a range of points of view about the way the world is, there is also a fairly clear way to resolve it, which is to design and perform another experiment. Things aren't so easy in respect of morality, though: there may be at least *prima facie* reasonable disagreement about what is good, and further reasonable disagreement about what the criteria for applying a word like 'good' are in the first place. In any given moral dispute, each participant will think his position correct – this much is implicit in his thinking that it's worth disagreeing about to begin with – but it is not a given that either is less rational than the other, or has failed in some aspect of moral perception. On the basis that it either is or isn't permissible to, say, kill a person at their request, it's impossible for both to be correct. Yet it's unclear whose moral positioning needs to be enhanced.

It's certainly not obvious that one of them is insufficiently empathetic and that boosting pro-sociality would make solving the dilemma any easier.

This is not hyperbolic scepticism: moral attitudes change over time and distance. What counts as hypocritical flattery in one place may be required by politeness in another; turning the other cheek may be a virtue for some, and seen by others as a blameable unwillingness to take revenge; bloodthirsty Aztec religious rituals were obviously not considered indefensible by the Aztecs who performed them. Aristotle has no time for a generalized philanthropy. We take it as read that justice is closely related to giving to people or punishing people according to what they deserve – but the *Odyssey* nevertheless speaks of Athene urging Odysseus

> to go round collecting scraps from the suitors and so learn to distinguish the good from the bad, though this did not mean that in the end she was to save a single one from destruction[36]

– which seems to show a flat contradiction of the principle coming from the mouth of a goddess: the goddess of justice, no less! It is possible that all these unfamiliar conventions are, as it happens, trading on a misapprehension of some sort – but until we can come up with a way to demonstrate that once and for all, there mightn't be all that much we can say to settle the matter.

Even within a single life – or within a single day of a single life – a disposition may be held to be desirable at some times, and utterly undesirable in others. To be disinterested and scrupulous about detail is a virtue in a judge, but becomes a vice when she removes her wig and goes home to her family. Were we to seek to enhance the moral motivations of judges, ought we to make them more disinterested? That would mean that we had made them vicious outside office hours.

A nice example of this trap is given in a video on moral enhancement on the website of Humanity+, formerly known as the World Transhumanist Association.[37] In this video (the transcript of which is available via his website[38]), Steve Omohundro talks of moral progress:

> If you look at recent history . . . there is a trend toward great moral progress. 200 years ago, slavery was generally accepted. Now, it is viewed as immoral almost everywhere in the world, at least officially, and pressure is put on societies that still allow it. The same is true of torture, though there has been a lot of recent controversy about it. We have the Geneva Convention and the notion of war crimes, the sense that war is bad but there are things within war that are especially bad. We have the establishment of women's rights in many countries, though some are still lagging. The same is true of racial equality. And the animal rights movement is growing rapidly.
>
> . . .
>
> When we think about transhumanism, I think we should start from humanitarianism. That is the notion that the things that most humans view today as precious, like human life, love, happiness, creativity, inspiration, self-realization,

> peace, animals, nature, joy, children, art, sexuality, poetry, sharing, caring, growth, contribution, spirituality, family, community, relationships, expression, are truly precious. These things matter because they matter to us. We may not know why these things matter to us, but that does not take away from the fact that they matter to us.

For sure: we have a set of social norms and conventions that we did not have a few centuries ago; and we (presumably) believe that the changes are for the better. But it's fallacious – and a touch arrogant – to suggest that 'the things that most humans view today as precious . . . are truly precious'. We cannot say, at least without further argument, that the reason why we prefer today's conventions is other than that they're the ones to which we're used – even allowing, extraordinarily enough, that our forbears had no time at all for any of the things he thinks 'truly precious'; this is a point nicely captured by John Mackie's suggestion that

> people approve of monogamy because they participate in a monogamous way of life rather than [participating in] a monogamous way of life because they approve of monogamy.[39]

To talk about moral progress in other terms is a bit like looking at Chaucerian English and talking about all the progress we've made in becoming better at spelling.

More particularly, it's hard to escape the worry that the progress mentioned in Omohundro's talk is all towards the kind of things that the wealthy, middle-class, educated people who tend to become transhumanists like. So when he continues to say that

> I think that the kind of morality and moral structures we want to create using new technologies should serve to preserve these qualities,

it's tempting to respond with a curt 'Well he would say that, wouldn't he?'. There is no suggestion that the moral characteristics of which these people approve may be merely the transient markings of a particular group in a particular epoch. He may have identified characteristics that are, as a matter of fact, desirable; but the problem is that there's room for serious doubt on that front.[40] (In fairness, we should remember that Omohundro has several companions in guilt: even Kant starts with the idea that his moral system will simply elucidate and take its lead from the 'common idea of duty and of moral laws'.[41])

The ARD applies equally to Persson and Savulescu's proposal – or any other attempt at dispositional engineering. In one sense, the desire that people should be more sympathetic and just[42] seems fairly trivial; but once we begin to put meat on the bones of what a just attitude would look like, we open the way for there to be reasonable disagreement. Persson and Savulescu take a particular understanding of justice as their touchstone, and think that we

should be time- and location-blind. But it is possible reasonably to think that one has commitments to some that one does not have to others; and, if this is right, then being dispassionate would be the *wrong* way to go.

The point of the ARD is this: it conflates giving someone better motives with giving someone what we *believe* to be better motives; and those beliefs may be false. This is admittedly not necessarily a fatal problem. It is possible to endorse a particular set of values while still being awake to the possibility that a radically different set is possible. In like manner, one may endorse a particular set of spellings while still being awake to the possibility that they will alter over the coming centuries. Hence, though for any two cultures *P* and *Q* with divergent sets of moral beliefs, members of *P* will believe things that the people of *Q* will think mistaken, people in either can quite coherently entertain a first-order belief about right and wrong at the same time as a second-order belief about the possible falsity of the first-order belief. Nevertheless, if we accept the formal moral desirability of giving someone morally better motives, we may be systematically mistaken in our attempts to realize our desires; and we may thereby unwittingly give someone motives that others may quite reasonably disavow.

Would-be moral enhancers may want to bite this bullet, and admit cheerfully enough that all they want, and could reasonably want, is to make people more likely to behave in ways that they prefer. But if this really is what's going on, then the word 'moral' seems to have been eviscerated – or more likely become equivalent to 'moralistic' – and we're back to the worries about freedom, dignity, and shaping others' ends for them. It's one thing to make predictions about what morally enhanced people would look like; another to make claims about the kind of people that we would like to see and call that moral enhancement.

For this reason, it may simply be foolhardy to attempt to instil particular dispositions in people without at the same time giving them the ability to reason morally and recalibrate those motivations should the situation require: we may have got ourselves into a muddle about what would be desirable.

## Enhancing moral reasoning

As we've seen, John Harris' attack on non-cognitive approaches to moral enhancement is substantially based on the idea that intervening to improve someone's dispositions makes them less free, and therefore less moral. For him, judgement is key:

> One thing we can say with confidence is that ethical expertise is not 'being better at being good', rather it is being better at knowing the good and understanding what is likely to conduce to the good.[43]

Elsewhere, he makes the much stronger claim that '[b]eing moral is like being scientific' and that the idea that emotionally-motivated behaviour is properly

considered moral 'is simply not credible'.[44] For Harris, morality is a cognitive process. Similarly, with Sarah Chan, he admits that we could use drugs to 'to treat anti-social personality disorders and social dysfunction'. But, he asks,

> would this constitute a 'moral enhancement', if the beneficial social effect were to be achieved through suppressing patients' own moral judgementeven if the outcome was a world with less harm?[45]

We might end up with 'nicer' people: but this may be moralistic enhancement rather than moral enhancement proper.

One thing that does seem to be clear in this part of the argument is that Harris' suspicion of non-cognitive moral enhancement such as that advocated by Douglas may not actually mean that many hits are landed. It's possible – and, I suspect, likely, that Harris and Douglas are talking past each other: the latter is talking about improving behaviour, and the former about improving thinking about behaviour (hence the point about moral *expertise*). As a result, it is not obvious to me that Harris' arguments really inflict much damage on Douglas'. Still, given the problems that face dispositional engineering as a strategy for moral enhancement, we might wonder whether some kind of cognitive moral enhancement programme could improve our moral reasoning and thereby generate more desirable behaviour.

It's a reasonable question, but it's not obvious that it would. One of the reasons for this is that improving a person's reasoning about what he should do is not necessarily going to be enough to make him do it – it will not remove weakness of will. Further enhancements to avoid this are a possibility – but not guaranteed to be desirable. More importantly, there's no obvious reason to think that morally undesirable behaviour really does have much to do with a failure of reasoning, so that being better at moral reasoning is going to improve behaviour; it might simply allow people more easily to convince themselves that their victims deserved to be treated poorly.

Imagine an antebellum Dixieland slave-owner; call him Zeke. Zeke is not an altogether bad man: he believes that one should not behave arbitrarily badly to one's slaves, and even thinks that he has duties to them: duties to provide them with food and basic medical care, to refrain from violence, and so on. However, Zeke also happens to believe that, whether it be by nature or divine ordinance, black people simply are inferior to whites, and properly (or at least not improperly) kept in bondage. The sympathy he shows to his slaves is, more or less, the same as the sympathy that he shows to his cattle: he does not allow his animals to starve, or to suffer illness or injury, or to be beaten, either. Still: he thinks that horses and cattle are not his moral equals, and that it is permissible to force them to work, sell them, and so on – and the same applies to his slaves. We'd be inclined to say, without much need for argument, that his beliefs about the status of his animals may be true, but the corresponding belief about his slaves is plain false; and this is at the root of his immoral behaviour in respect of them.

This example might be taken to speak in favour of a cognitive account of moral improvement. Harris articulates something like this thought:

> The most obvious countermeasure to false beliefs and prejudices is a combination of rationality and education, possibly assisted by various other forms of cognitive enhancement, in addition to courses or sources of education and logic.[46]

It is certainly not a given that Zeke's moral failure would not be remedied by making him more empathetic or pro-social: his problem is not so much a lack of empathy, as a mistake about how it should be directed. But for his beliefs about race, he would be an uncontroversially good (or, at worst, much better) person: and so if we want to make Zeke more virtuous, all we have to do is shift that belief about the moral worth of a portion of humanity. It's not clear that the biosciences could help us here – and so in terms of my wider investigation about the contribution that those sciences might make to the good life, we could turn out to be on a hiding to nothing. But even ignoring that, the problem for the cognitive account is that it is not obvious that that would generate a moral improvement anyway.

A significant part of this is that, if Zeke's underlying assumptions about his slaves are that they are at least somewhat lower than himself in terms of moral importance, and there *is* a fault in his moral reasoning, we need to distinguish between correcting the faulty reasoning, and correcting the false belief on which it builds. Imagine that Zeke teaches some of his slaves to read; correcting Zeke's reasoning might lead him to the conclusion that he is thereby wasting good money treating them better than they really warrant, given their inferior moral status. Harris has faith that making people better at moral reasoning will make them more likely to behave in a morally desirable way; but it's at least possible that it'll have the opposite effect. There seems to be no guarantee from the start that someone like Zeke would, by thinking more clearly about how to treat black people, decide that he ought to treat them better. The problem is with the underlying false belief about the moral status of certain humans – but this might not be the kind of thing that can easily be corrected by cognitive sharpening: falsity is not a sign of irrationality. (The same criticism applies to Persson and Savulescu's plea for generational blindness: however admirable such an attitude, it's not obvious that it's particularly *rational*.)

Indeed, there's reason to be fairly confident that improving someone's reasoning would not necessarily make a person's moral behaviour better, because behaviour is often informed by inbuilt prejudice, and that is not always open to reform by clearer thinking. Such prejudice – be it against members of a particular ethnic group, or one sex, or people with certain disabilities, or whatever – can persist even when a person honestly believes that it is not there; it can persist even when there are good reasons to expect that it

would have no cognitive traction at all. A good example of this is provided by the American political activist Jesse Jackson, who is widely quoted as having said that

> there is nothing more painful to me at this stage in my life than to walk down the street and hear footsteps and start thinking about robbery, then look around and see somebody white and feel relieved.[47]

This statement claim is important in respect of the idea that cognitive enhancement can eliminate the basis for undesirable behaviour such as explicit racism, because the idea that racial prejudice is reducible to false beliefs, which would be eliminable by clearer thinking – and that racist behaviour would thereby never appear – seems to wither under its light. After all, an educated black civil-rights activist is rather more likely positively to disavow than hold a conscious belief that black people are more dangerous; and yet a prejudice against black people seems to linger. I'm not, of course, suggesting that Jackson has ever behaved in a racist way; but if even his underlying attitudes can be influenced by underlying and consciously disavowed prejudices, it would not be surprising to find that at least some racist behaviour in others is influenced in the same sort of way. Eliminating false beliefs is not enough to improve behaviour.

This is instructive in respect of Zeke. It might be that he has a set of conscious beliefs in respect of black people; and it might be that education (in the sense of learning to think more clearly: the sense that Harris seems to mean in the quotation above) or some other form of cognitive enhancement, would alter them for the better. Indeed, the idea that one can draw some of the venom from undesirable moral attitudes by means of education is intuitively plausible. By such strategies, we might well be able to get people to take on board that their attitudes and behaviours are ill-founded and indefensible. But there persists the worry that such interventions would not be enough: if Jesse Jackson is vulnerable to a kind of cultural conditioning about black people that persists despite having in his own use of the first-personal pronoun plentiful evidence that black people are not particularly threatening, then so would be Zeke; and getting Zeke genuinely to believe in the moral equality of black and white people would not obviously be enough to shift the underlying prejudice if it's not enough for someone like Jackson. This, incidentally, gives us another brickbat to throw at the idea that moral failure is a failure of reasoning: increasing Zeke's moral reasoning abilities might just mean that he is better at coming up for a *post hoc* rationalization of his unexamined-because-unconsciously-held attitudes.

In short, if it turns out that racism has, or can have, its roots in some non-cognitive characteristic, cognitive enhancement is unlikely to be enough to get rid of it. There is, then, at least one situation in which improved cognition will

not lead reliably to improved disposition. I suspect that this is enough to break the back of the idea that improving moral reasoning will make people morally better, and that John Harris' claim that

> if [love, altruism and the like] are not felt, or worse not felt for particular groups, racial, national, religious political, or for particular individuals, that is when you need moral philosophy, moral reasoning[48]

is, on this ground, open to a charge of naïvety.

By the same token, if certain undesirable attitudes worm their way into our souls on account of our socialization, then unless the social *milieu* changes, these attitudes will stand a chance of remaining; but, more optimistically, the hope remains that a change in the social *milieu* could be all that's needed to eliminate them. Again, bioscientifically informed procedures may well be *de trop*. Moral enhancement in the broadest sense might still be something that we think worth achieving – but it would be more-or-less of the same stripe as the kind of common courtesy that we instil as a part of growing up. It does not generally require oxytocin or giving kids particular beliefs: it's a matter of acculturation.

And the same could well apply in respect of making people morally better in the round: using bioscientific interventions to improve problematic behaviour could well be, at the very least, like using a monkey-wrench to crack a monkey-nut. Once again, we're back to the idea that it's habituation that might be able to carry most, if not all, of the weight of making people better. Once again, Aristotle seems ahead of us:

> Moral goodness is the result of habit, from which it has actually got its name, being only a slight modification of the word *'ethos*.[49]

It's not beyond the realms of possibility that that's where the most progress is to be made in respect of moral enhancement.

## Is moral enhancement desirable anyway?

If proponents of dispositional engineering and cognitive moral enhancement are talking past each other, this might be a sign that their positions are not, in the end, incompatible. I suspect that this might be correct – though, if the ARD is well-targeted, we might do well to hesitate before trying to combine them. However, it's also worth noting another aspect of the ARD that has not been mentioned before. Should we make a mistake about the best dispositions to instil in others, and the things we thought desirable at time *t* turn out to be not desirable at *t+1*, this would be bad, or at least non-optimal, from our point of view. But what about the point of view of the enhancee?

Go back for a moment to the Frankenstein example, where Dr Frankenstein gave his monster a set of desirable dispositions. All being well, the monster's dispositions and his moral thinking would cohere. But if the ARD is correct, this might not be guaranteed: the monster may decide that the dispositions that he has had instilled are mistaken; and this would seem to bring a significant risk that he would be subject to a series of dispositions that he morally disavows. This may well undermine his quality of life: he would be alienated from his own motivations, and such alienation may be distressing.

What goes for monsters goes for humans: to have the freedom to identify a cause, but to be caused to act otherwise, is the kind of thing that we may reasonably expect to be incompatible with a good life; and if the point of the biosciences (and enhancement in particular) is at least partially to make people's lives better, such an eventuality would appear to undermine the point *ab initio*.

Correlatively, though one might expect all people to be morally good in the broad sense, this is not enough for them to count as leading a good life. As Ottfried Höffe points out, 'a morally good life is no sufficient condition for a totally good life'.[50] For him, though '[t]he morally good life is a condition for the art of living',[51] and though there may be a tension between the art of living and the moral, and though even 'the moral takes precedence', still, it does so

> not so exclusively that it eliminates all rights to the art of living. . . . As a consequence, a philosophical ethics is not satisfied with a philosophy of morals. It recognises happiness as an independent, admittedly subordinate principle.[52]

The idea that the moral takes precedence over the art of living is not something we should accept as a given; but the final sentence of the quotation does capture an important point. Moral enhancement pursued as an end in its own right ignores the fact that moral goodness is at most a *part* of a life worth living. Moral enhancees' lives have not necessarily been enhanced in the round. If we're interested in the promotion of the good life, that matters.

In other respects, moral enhancement may not be the best solution to a problem faced by a would-be enhanced person. Take, for example, the idea that we could intervene pre-natally to optimise a person's MAOA functioning. Savulescu et al. note that it is not only the people around someone who is neuro-genetically disposed to antisocial behaviour that suffer:

> Given that compulsive violent or criminal behaviours are typically punished with incarceration as well as social rejection, we can assume that these behaviours typically harm a person who expresses them, as well as their victims. Selections which constrain a person from making such harmful choices may therefore be justified.[53]

Still, while there may be harms that accrue from being at the receiving end of the criminal justice system, this could just as easily be taken as a call for the reform of that system as for moral enhancement. Indeed, Savulescu et al. admit the possibility that

> [a] deficiency of MAOA results in a build up of neurotransmitters. Once they reach abnormal levels, this is thought to *increase the probability* of a person demonstrating an excessive, even violent, *reaction to stress* – a theory demonstrated in knockout mice.[54]

The emphasis is mine, and it seems important. One reason is that genetic cues tell us about probabilities rather than certainties. A second, and more important reason, is that if the MAOA hypothesis is true, a stressful environment is at least a companion in guilt when it comes to causing antisocial behaviour; people with the mutation in question might benefit from having access to (for example) a pill that would help prevent stress; but even here, this would – it seems to me – count as therapeutic rather than enhancing, and not necessarily *morally* therapeutic at that. And, of course, reducing socio-economic stressors would also, presumably, improve the lives of others living in the same circumstances who have the good fortune not to have abnormal MAO activity.

That there is a possible role for environmental factors in understanding and ameliorating undesirable behaviour is no small point. There is evidence that links antisocial behaviour with exposure to lead pollution, for example.[55] This may well serve as an indication that to focus on moral enhancement as a solution to social problems is to misidentify the root of those problems: if people's social development is retarded or skewed by exposure to pollutants of one sort or another, they do not need bare enhancement; the need therapy. More importantly, it may be possible to improve behavioural patterns without having to worry too much about manipulating individuals' brains, and so enhancing technologies in the sense understood by many of their proponents may be *de trop* (although, admittedly, this would likely have more of an impact on subsequent generations than on anyone who currently finds himself with a brain affected by pollution). Finally, everyone would benefit from a cleaner environment anyway, irrespective of whether or not they had an elevated propensity for criminal activity: Elise Gould estimates that

> [f]or every dollar spent on controlling lead hazards, $17–$221 would be returned in health benefits, increased IQ, higher lifetime earnings, tax revenue, reduced spending on special education, and reduced criminal activity.[56]

Behaviour isn't the only aspect of life that'd get better.

Finally, there is no guarantee that moral enhancement would be good for humanity as a whole: there could be vices that are worth keeping. If that is too much of an oxymoron, we could still talk about characteristics that, while undesirable when manifested in individuals, are indirectly beneficial.

In his *Fable of the Bees*, Bernard Mandeville prefigures Adam Smith by voicing the idea that the common wealth may be increased by the action of self-interested behaviour: 'Bare Virtue can't make Nations live/In Splendor',[57] and there may be public virtue grown from the soil of private vice. This

sort of claim might be spun out thus: that a great deal of economic activity is generated by people wanting to maximize their own welfare, and would not take place if people were altruistic, or inclined to treat others' welfare as being as important as their own. Such activity, over time, generates welfare for everyone. In reality, things are much more complicated: though there may be a 'trickledown effect' from private vice, there may be problems related to market failure, widened inequality; and the benefits may take years to dissipate, which is of little comfort to those living in need right now. Still: a good deal of current economic theory tells us – insists, in fact – that, overall, the most stable and efficient way to increase general welfare is to allow people *not* to be altruistic. The trick of the just society is to harness those vices and channel them rather than eliminate them:

> So Vice is beneficial found,
> When it's by Justice lopt and bound.[58]

The lopping and binding are crucial: but this is not a reason to opt for bioscientific moral enhancement. What actually counts as the just distribution of resources is a matter for reasoned public debate; getting people to behave in a way that justice dictates is what public policy is for.

This is a pattern that could be repeated more or less indefinitely. Indolence may be a vice, but a desire for ease can, if appropriately harnessed also be the spur for labour-saving devices that make life better, for example. Elijah Millgram's essay comparing Hume and Gibbon in their accounts of the fall of Rome points out that – bluntly – Gibbon thought that Rome fell because it became too immoral while Hume thought that it fell for just the opposite reason: too much morality.[59] On the assumption that Hume thought the fall of Rome to be a loss, this would appear to generate the idea that too much morality – or too much morality of a certain kind – might be bad over all. There's something in this reminiscent of Orson Welles' 'cuckoo clock' speech in *The Third Man*:

> After all, it's not that awful - you know what the fellow said . . . In Italy, for thirty years under the Borgias, they had warfare, terror, murder, bloodshed – they produced Michelangelo, Leonardo da Vinci and the Renaissance. In Switzerland they had brotherly love, five hundred years of democracy and peace – and what did that produce . . .? The cuckoo clock.[60]

Historically inaccurate as it may be, the speech captures nicely the idea that not all admirable things are godly to the core. There is a sense in which wholly 'moral' lives within a wholly 'moral' world would be lives that noone would want to live, in a world that no one would want to inhabit.

This is not a defence of vice *per se*: it might well be that there are all kinds of traits that are properly listed as vices that we would rather be without,

and others that are properly listed as virtues that we would rather see more frequently, and that no reasonable people would disagree. But there might be indirect benefits to be had from retaining the possibility of at least some other vices; and – more problematically – it might be that certain undesirable behaviours are coeval with other, desirable, ones.

I would want to be among the first to admit that these qualms about moral enhancement are not definitive. However, I think that they are enough to start to pick away at the idea that moral behaviour can easily be altered for the better by bioscientific means. The standard by which improvement is to be measured is unclear; the benefits to society are unclear; and the enhanced individual, for a range of reasons, might have defensible objections to being enhanced; and, most importantly, even when moral enhancement is possible and desirable, there really isn't any reason to suppose that bioscientific interventions are the wisest method of pursuing it.

Moral enhancement by habituation and education leaves room for nuance, for judgement, and for altering not just the person, but the interplay between him and his environment. Bioscientific interventions, it seems to me, are much too rough-and-ready, and take too much for granted.

## Notes

**1** Douglas, 2008, p. 229.
**2** Ibid., p. 231.
**3** Ibid.
**4** Savulescu et al. 2006, *passim*.
**5** Persson and Savulescu, 2012, *passim*.
**6** Ibid., p. 107.
**7** Savulescu et al. 2006, p. 165.
**8** Harris, 2011, p. 110.
**9** Ibid., p. 104.
**10** Ibid.
**11** Kant, 1993, 4:398–9.
**12** Savulescu et al. 2006, p. 169.
**13** Fine, 2011, *passim*, but esp. chs. 9–11.
**14** Persson and Savulescu use the word 'causally' rather than 'neurophysically'; I've made this substitution because of the use that I'll be making of the word 'cause' later in the chapter. It doesn't detract from the sense of the argument, though.
**15** Persson and Savulescu, p. 112.
**16** 'Persons, Character and Morality', in Williams, 1999, p. 18.
**17** cf. Persson and Savulescu, *op cit*, p. 113.
**18** 'Are Moral Requirements Hypothetical Imperatives?', in McDowell, 1998, p. 85.

**19** Ibid., pp. 90–1.
**20** DeGrazia, 2013.
**21** Ibid.
**22** Harris and Holm, 2003, pp. 125–6.
**23** see Lyons, 2012, pp. 371–2.
**24** Aristotle, 1976, p. 91.
**25** Ibid., pp. 91–2.
**26** Ibid., p. 92. It's interesting to note that Aristotle associates ethical formation with the *polis* rather than the home. Maybe he thinks that the importance of the domestic is self-evident – except that elsewhere in his writings he has no compunction about weaving commonplaces explicitly into his investigations.
**27** Aristotle, 1976, p. 92.
**28** see Anon, 2009b and Anon, 2010b.
**29** Douglas, 2008, p. 229.
**30** Ibid.
**31** Harris, 2013.
**32** *contra* Shook, 2012, p. 4; cf. DeGrazia, 2013.
**33** Savulescu et al. 2006, *passim*.
**34** see Harris and Chan, 2010, p. E183; Chan and Harris, 2011, p. 131; Harris, 2011, p. 106.
**35** Douglas, 2008, p. 231.
**36** Homer, 1991, VXII: 360–4.
**37** http://www.humanityplus.org/2009/05/steve-omohundro-on-ai-and-the-future-of-human-morality/.
**38** Omohundro, 2009.
**39** Mackie, 1977, p. 36.
**40** In a similar vein, when Agar (2013) expresses a concern that post-humans might have a higher moral status than mere humans, and come to see mere humans as things that could be sacrificed in the pursuit of post-human interests, Douglas (2013) is confident in his response that the transition to post-humanity 'would, after all, make them less liable to sacrifice' (p. 76). The thought here seems to be that enhanced post-humans would express certain moral traits as a matter of course, and that these traits are predictable. I see no reason why we should be forced to accept this – though I see no particular merit in Agar's worries, either.
**41** Kant, 1993, 4:389.
**42** Persson and Savulescu, 2012, p. 108.
**43** Harris, 2010, p. 104.
**44** Harris, 2012, p. 297.
**45** Chan and Harris, 2011, p. 131.
**46** Harris, 2010, p. 105.
**47** The source of this quotation appears to be an article in the *Chicago Sun-Times* from the 29 November 1993 called 'Crime: New Frontier – Jesse Jackson Calls It Top Civil-Rights Issue' by Mary A. Johnson. I'm grateful to a pseudonymous comment on the *Feminist Philosophers* blog for this information; however, I have not been able to get access to the original to verify it.

48 Harris, 2012, p. 3; slightly modified.
49 Aristotle, 1976, p. 91.
50 Höffe, 2010, p. 331.
51 Ibid; slightly modified.
52 Ibid, p. 333.
53 Savulescu et al. 2006, p. 165.
54 Ibid., p. 159.
55 see Nevin, 2007; Mielke and Zahran, 2012.
56 Gould, 2009.
57 Mandeville, 1997, p. 35.
58 Ibid.
59 Millgram, 2005, ch. 8, *passim*.
60 Greene and Reed, 1968, p. 114.

# 8

# Bioscience and the Duty to Research, Part 1: Ways to Make Life Better

If we cast our minds back to the end of Chapter 3, two claims should have been apparent: first, that there was no duty not to enhance ourselves or others; second, that there was no duty to do so either. This leaves the possibility that there are still reasons to pursue enhancement for the sake of the good life – but the last four chapters have been devoted to the defence of the claim that those reasons are not particularly compelling, and that interventions to promote the good life need not be bioscientifically-led. The insights generated by the biosciences are most likely to make a difference for the good when they're put to work in the pursuit of ostensibly trivial things like sport and art. Generally speaking, there is no pressing reason to pursue bioscientific innovation down the road towards enhancement.

But this isn't going to be enough to tell us that there is no reason to pursue bioscientific innovation; and there *is* a reason to do so. Concentrating on what I called 'bare enhancement' makes it easy to forget that therapy is a kind of enhancement; and there are conditions that generate reasons to pursue bioscientific research for therapeutic reasons – that is, for the sake of being able to guarantee at least a minimally acceptable quality of life for their beneficiaries. I'm willing to accept more or less as axiomatic that the recognition that the quality of a life falls below a level we would find acceptable generates a reason to do something to improve it. Such improvements may be generated by means of research in the biosciences; hence we may have a reason to pursue such research. The question that I want to address in this chapter is this: is that reason overriding, amounting to a duty – incumbent on individuals, or on society as a whole – to pursue bioscientific research?

I'll be claiming that there is not a duty to research for the sake of promoting the good life. In the next, I'll look at the idea that there might be a duty to research based on other considerations, and claim that the case for such a duty has not been compellingly made there, either.

The first question we have to ask is whether anyone thinks that research in the biosciences is a duty to begin with. Certainly they do. John Harris (latterly, alongside Sarah Chan) has claimed explicitly that participation in, or at least support for, scientific research – particularly research in medical and bioscience – is a moral duty. That one of his papers is called 'Scientific Research is a Moral

Duty'[1] is ample to establish this. (From hereon, I'll use the phrase 'duty to research' for the sake of ease.) This is a claim that, he thinks, follows from a much broader claim to the effect that '[t]here is a universal responsibility to pursue the good, for ourselves and others'.[2] It's quite clear from the context that 'responsibility' here should be taken as sufficient to indicate a duty. If it's correct to say that there's a duty to pursue the good, if research is a means of doing so, and if nothing is more immediately demanding of our attention, then research could count as a way of discharging the duty; and if it's the *best* way of pursuing the good, it would be obligatory. And though Harris will attract the lion's share of my attention over the coming pages, plenty of other people argue that there's a duty to research, as we shall see over the coming pages.

As in Chapter 3, I'm going to treat the word 'duty' as meaning that action the omission of which would reasonably be blameable. I'm also going to assume that the duty is owed to at least one other person or thing; for while there could well be self-regarding virtues and duties, and that pursuing one's own good might coherently be one, it would be unclear who would be bothered if one fell short of discharging such a duty. There might be all kinds of reasons and incentives to pursue the good on your own behalf, but to talk about duties and responsibilities in this context suggests more than that. (Treating a reason to act in one's interests as a duty is a trap into which Stjernschantz Forsberg et al. fall.[3]) A responsibility to pursue the good is something I take to make most sense when it's at least related to a duty of beneficence. And, indeed, an appeal to beneficence might well be precisely the sort of thing that'd plausibly provide a foundation for a duty to research anyway.

Other things hinge on what we mean by 'pursuing the good'; as we'll see in what follows, even if we can get everyone to agree that we have a responsibility to pursue the good, there could be reasonable disagreement about what this entails. Obviously, given the interest that this book has in 'the good life', this is an important question. All in all, *any* claim that there's a duty – or, more softly, a responsibility – to pursue scientific research has at the very least a number of notable questions it needs to confront. The hypothesis that I want to advance is that it can't answer them.

## Is there a duty of beneficence?

It's very tempting to try to ground the idea that there's a duty to research in some kind of appeal to beneficence: that in pursuing research, we are contributing to an improvement in the world and the lives of the people therein, and that it is that contribution that makes research morally valuable. But something's being valuable is not the same as it being obligatory; and if we are sceptical about whether beneficence is obligatory, then we can make all the noise we like about the link between it and research: it won't after all tell us that participation

in, or support for, research is required. I can't begin to offer a comprehensive analysis of the questions that surround a proposed duty of beneficence here – much less to answer them. But I can begin to add a little to the sketch of reasons to be sceptical about such a duty that I offered in Chapter 3.

I shall not deny that it is good and admirable to be beneficent. I'll also take as read that we have a *ceteris paribus* duty to avoid causing harm. Nevertheless, a duty not to make the world a worse place is not the same as a duty to make it any better, just as my duty to rob you is not interchangeable with a duty to make you wealthier; and establishing that there's a positive duty to confer benefit other than by *fiat* is not all that easy a task. How might the idea that there's some kind of duty of beneficence be bulked out?

One strategy might be to invoke Mill. Mill thought that he had demonstrated in *Utilitarianism* that happiness – which we may construe more widely as welfare – was the measure of the good, and the only thing desirable as an end ('there is in reality nothing desired except happiness'[4]), and that we ought to pursue the greatest happiness for the greatest number. Yet the problems with Mill's attempt to aggregate are famous. In the first place, he manages to fall victim to his own sleight of hand by moving from a claim about happiness being desirable in the sense of capable of being desired, to a claim about it being desirable in the sense that we ought to desire it. Hence he is vulnerable to the charge that he begs the question, being able to argue that providing happiness is a duty only inasmuch as that he has helped himself to this conclusion already.

The problem is one of translating good into duty – of making the leap from a claim about what does motivate people's actions (primarily in respect of themselves), and what ought to motivate people's actions (especially when the good supposedly desired accrues to others). Attempts to ground obligation in empirical claims about motivation are not up to the job. The obvious response is simply to bite the bullet, admit that there's a gap between what we're likely to do or what rational self-interest tells us we ought to do for prudential reasons, and what we ought morally to do, and say that the job of public policy is to bridge it. If we acknowledge that research is welfare-maximizing and that we ought to maximize welfare, then policymakers would have a warrant to generate the kinds of carrots and sticks that are necessary to get people actually to support research. This response is fair enough – but it still takes as read the idea, which seems to be central to the claim that we should pursue research as a means of beneficence, that we ought to *promote* welfare, rather than simply recognize it as a good thing.

That it should be promoted is something else that Mill thinks he can demonstrate. His claim in this context is that, since his own happiness is desirable to each, so the aggregate of people will desire the good for the generality.[5] But this misses two things. The first is that the aggregation of individual agents desiring their own happiness will not – certainly not with any certainty – give us anything except a fractious crowd of people who each

desire their own individual happiness. What is in the interests of everyone is not the same as what is in the interests of each; and in trying to make the jump from the claims about each to claims about all, we need to take care to ensure that the desires of each are mutually compatible. In the context of claims about a duty to research, this is a serious problem, because each of us might have different accounts of what is necessary for happiness; it is far from certain that what I would do for the sake of my happiness is what you would do for the sake of yours.

The second is that, even if we could aggregate, it's unlikely that an individual would be indifferent between his actions generating happiness for himself and those same actions generating happiness for the generality. If community-mindedness does not generate as much happiness for the individual as would selfishness, then whatever reasons there might be for community-minded action would not be derived from the same source as those for self-interested action – which, again, means that the aggregation claim is unconvincing.

Again, there is a problem in moving from incentives to duties; but in this case, it's combined with a problem about scaling up from individuals. Hence there is no guarantee that we would be able to come up with any coherent plan about what research to do, even if we were satisfied that research of some sort ought to be on the agenda.

Once again, we might say that it's the job of public policy to fill these gaps; but that assumes that they *should* be filled, and this presupposes that there is a duty to be other-regarding; and so the important moral question remains open. And even if public policy was designed to increase the likelihood of us performing actions that benefit the generality, this position all but amounts to abandoning beneficence properly so-called: what gets people to contribute to the public good on this account is not the fact that it is the public good, but the straightforward fact that policy has been designed to incentivize certain desired behaviours. The most that can be said is that we end up with what generalized beneficence would have given us if it gave us anything. There is nothing new about this argument; but it's worth noting from the point of view of conversations about a beneficence-based duty to research.

Another attempt to ground a duty to research in beneficence, this time based in deontological thought, offers to be able to fill these gaps more successfully. As with utilitarians, deontologists have to provide an argument that will show that putative duties really are duties. The *locus classicus* when trying to establish a duty of beneficence seems to me to be the argument that Kant presents about a duty to aid as set out in the *Grounding* and the *Metaphysics of Morals* – an argument that I mentioned in Chapter 3. At least by Kant's normal standards, the argument in both texts is pretty clear. His opening gambit for establishing the duty in the *Grounding* is to ask us to imagine a man for whom all is going well; he sees others whom he could help 'struggling with great hardships', but decides not to assist:

> Let everybody be as happy as Heaven wills or as he can make himself [our protagonist says]; I shall take nothing from him nor even envy him; but I have no desire to contribute anything to his well-being or to his assistance when in need.

Kant admits that this maxim of indifference could exist as a law of nature. But, he adds,

> still it is impossible to will that such a principle should hold everywhere as a law of nature. For a will which resolved in this way would contradict itself, inasmuch as cases might often arise in which one would have need of the love and sympathy of others and in which he would deprive himself, by such a law of nature springing from his own will, of the hope of aid he wants for himself.[6]

Much the same argument is presented in *The Metaphysics of Morals*:

> To be beneficent, that is, to promote according to one's means the happiness of others in need, without hoping for something in return, is everyone's duty.
>
> For everyone who finds himself in need wishes to be helped by others. But if he lets his maxim of being unwilling to assist others in turn when they are in need become public, what is, makes this a universal permissive law, then everyone would likewise deny him assistance when he himself is in need, or at least would be authorized to deny it. Hence the maxim of self-interest would conflict with itself if it were made a universal law, that is, it is contrary to duty. Consequently the maxim of common interest, of beneficence towards those in need, is a universal duty of human beings.[7]

Beneficence, the promotion according to one's means the happiness of others in need, is only an imperfect duty – that is, it is impossible to *will* coherently that it should be a law, but not logically impossible for it to be one – but it is still a duty for all that.

I argued when I first considered this argument that it does not establish that there is a duty to enhance; does it fare any better in demonstrating that there is a duty to pursue beneficent research? Probably not. This is for a couple of reasons. The most important is that it seems pretty clear that the word 'need' is important: Kant is setting out an argument against stinginess in the face of situations in which people's suffering is an immediate concern. He does not appear to endorse any stronger duty about generalized beneficence in the absence of particular need. As a matter of wringing the most rhetorical power from his argument, had he anything stronger in mind, he would have been well advised not to appeal simply to situations of need when arguing for the duty.

The second reason is that, even if a wider claim is intended, along the lines that there is a generalized duty of beneficence irrespective of need, it wouldn't be convincing. This is because it is *always* possible to do good for others; hence treating beneficence as a universal duty in this sense would seem to require beneficence at all times, which amounts in practice to treating oneself as a means to realizing those others' good. Not only would this itself violate the

Categorical Imperative – we must, after all, treat humanity not only as a means but also as an end in itself regardless of whether that humanity is in ourselves or others – but it would also be vulnerable to a slightly reworked version of the 'integrity objection' to utilitarianism articulated by Williams,[8] inasmuch as that it would require that we see people's lives as being given over to providing benefit to others.

From this limited claim about a duty of beneficence, we are not required to admit that love and sympathy need amount to anything much more than acting to ensure minimal decency. And what this means for the sake of our argument here is that the fact that research stands a chance of making the world a better place than it might be will not, in and of itself, show that we have a duty to research, because Kant seems only to be talking about a duty of beneficence in cases need – and being in a position to benefit from something is not the same as being in need of it. However, research probably isn't mandated in times of acute need, as I'll argue in the next chapter. And, of course, Kant's demand can only plausibly be that we assist people according to our means: if – say – the cure for cancer does not exist, this means that our obligations will be manifested in terms of an obligation to provide what medication we can, or to provide non-medical support. It doesn't so clearly indicate that we need to expand the means available to us.

More importantly, even in these situations of acute need, it does not seem too implausible to imagine someone who would expect and want nothing more than *care* from others, rather than anything more radical. While it's likely that most cancer sufferers would welcome more research into a cure being done, this desire is compatible with seeking other manifestations of beneficence at least as much; and it's possible that bioscientific innovation is welcome, but not necessarily what is morally required. I'll make more of this idea later in the chapter.

(A certain kind of Kantian retort might be offered at this point. Kant claims in the *Grounding* that 'the ends of any subject who is an end in himself must as far as possible be my ends also, if that conception of an end in itself is to have its full effect in me'[9]. Doesn't this mean that, if a person decides that (say) living longer is his end, then I must adopt his longevity as my end also? And if that requires research into new drugs, then isn't there at least an obligation grounded in a hypothetical imperative to pursue or support such research, such that *if* longevity requires research, *then* we ought to research? Actually, no: this argument is naïve. The mere fact that something is someone's end can't *de facto* mean that I have obligations to realize it – or at least not overriding ones – on pain of not treating the humanity in me as an end. We have to ask what a reasonable attitude to that end would be, and how far it ought to inform our activity. And on the basis of the argument in Chapter 4, it would appear that life-extension is not automatically all that important a consideration. The same sort of consideration applies in respect of other ends.)

This point probably demands a lot more attention than I can give it here; but the point I want to make is that, overall, establishing a claim that there's a duty to research based on a duty of beneficence is not a straightforward task, because establishing that beneficence is mandated by duty is not easy.

## Beneficence, benefit and obligation

But let's assume, for the sake of the argument, that some kind of duty of beneficence can be established: will this suffice to tell us that *research* – particularly bioscientific and biomedical research – is a duty? Not necessarily.

The first thing to note is that there might be a point at which actions that satisfy a strict definition of beneficence do not really match with what a duty of beneficence implies on a more intuitive level. Imagine that there exists a wealthy industrialist who sets up a trust fund specifically to pay the university fees of the sons of millionaires. We would be hard-pressed to deny that this is beneficent. But this munificence will probably not make all that much of a difference either to the lives of the recipients or anyone else; if we're talking about a duty of beneficence, this would probably not be quite the sort of thing that we had in mind. What seems to be more in line with our intuitions about what a duty of beneficence would look like would be something more general – something like an obligation to make the world a better place in a more readily perceptible way.

Already, this speaks to what a putative duty to research would look like. Since there are several million people around the world whose lives do not meet what we would accept as a reasonable criterion of goodness, it would appear that they would have first call on our attention when it comes to beneficence: acting to improve the lot of those whose lives are already at least minimally decent seems to be a much less pressing requirement, and is probably supererogatory rather than required. However, it is not at all clear whether further research is quite what the worst-off people need. For while a malaria cure would improve lives and would be welcome, the reasonable response to the suffering of those in need might turn out to be that we should offer what relief we can. And (as I mentioned a little while ago) a duty to take what means we can to improve lives does not imply a duty to increase the scope or number of those means. A person's life may be sub-optimal, but still at least reasonably decent, if we can provide decent treatment when the fever does rise – not least because 'reasonable' will be informed by what's currently available. (Of course, if someone does come up with an effective, cheap and simple cure, our duties might alter to encompass providing that.) Correspondingly, if we think that there's a duty to prevent infection in the first place, that might be manifested as a duty to do what we can to drain the swamps in which Anopheles mosquitoes breed, rather than necessarily researching a vaccine.

And even if, implausibly enough, there is a situation in which research *can* be shown to be the only thing that fits the moral bill, this will tell us at most that there is a duty to support *certain kinds* of research, rather than research *in toto*.

For the sake of clarity, it's worth iterating that, in accordance with the model sketched out in Chapter 1, it might well still be the case that there would be all manner of innovations – perhaps motivated by market stimuli – and that gradually we would come to expect that different standards of acceptability be met: today's desirable life might be tomorrow's minimally acceptable one. Such eventualities might generate new duties. However, unless there is a duty to innovate to the extent that what counts as acceptable today may not tomorrow – an idea that is barely credible – then there doesn't seem to be all that much more to be said. No duty would be violated if, the lives of all having successfully been made at least minimally acceptable by the standards of the day, the economy and culture generally entered a steady state henceforward: as Hans Jonas wrote in 1969:

> Unless the present state is intolerable, the melioristic goal is in a sense gratuitous, and not only from the vantage point of the present. Our descendants have a right to be left an unplundered planet; they do not have a right to new miracle cures.[10]

So a (possibly narrow) range of fairly specific duties might be the kind of thing that we could ground in a duty of beneficence; while doing more than this might be admirable, and the kind of thing that more-than-minimally-decent people might well at least consider, it is not obligatory. It is not obvious that research falls into this narrow range. By contrast, if there's an ongoing duty to research, and if beneficence has anything to do with grounding it, then – at least so far – it would seem that we'd have to accept that it doesn't necessarily have to do so much with providing benefit in response to need, as with providing benefit *per se*. And for reasons that I outlined above, that might not always be that compelling a candidate for inclusion on the list of duties.

## What would be beneficial research?

But, notwithstanding the last paragraph, let's allow that the putative duty to research is wide enough in scope to cover the provision of benefit to those not in acute need. There's still a range of questions that warrant being asked. Notably, we might reasonably begin to wonder who the beneficiaries of our research ought to be, and whether the kind of benefit matters. As I write this, for example, NASA's *Curiosity* rover has already spent several weeks trundling across the surface of Mars, zapping bits of rock with a laser, and analysing the gasses released as it goes. This mission promises to tell us a great deal about

the history of Mars, some of which might help us work out if there could have been life there, and to make contributions to planetary science generally. It is, unquestionably, research. There would seem to be plenty of scope, though, for saying that one thing that it is not is beneficent – certainly not in the sense of being the kind of thing that gets a person a reputation as a philanthropist.

Yet this does not mean that it would be appropriate to call the mission a waste of time, effort, and money. Arguably, it does contribute in an important way to human wellbeing, and so could plausibly be counted as beneficial. This is for a couple of related reasons. The first is that missions like this promise to provide questions to answers that at least someone – and probably a good many more than one person – finds worth asking. Some value is attributed to finding an answer; but the attribution of value to its outcome is the only thing that it would make sense to think makes any research whatsoever worth doing. This does not preclude scepticism about whether the money given to NASA might be better spent elsewhere – but I'll come back to this problem in a little while.

But I think we can push this a little further, and make the suggestion that, even if people do not have a particular hankering for the answer to a given question, the kind of thing that *Curiosity* is doing is worthwhile *just inasmuch as that it makes lives richer*. The same sort of point could be made in respect of any number of other avenues of research. The discovery that the feathered dinosaur *Sinosauropteryx*, which lived in the early Cretaceous, probably had dark chestnut-coloured stripes[11] is perhaps not the most useful piece of information. However: this kind of discovery is – to use a term not used in academic work nearly often enough – *very cool indeed*. It is, in a non-trivial sense, life-enhancing. We are richer for it, and for many other useless discoveries.

This is no small point. Even if we think that duties of rescue mean that the primary obligations we face have to do with ensuring that a certain threshold of basic need satisfaction is reached, once it *is* reached, whatever duties to research remain – still allowing, *arguendo*, than any remain at all – will derive their merit from their ability to make an already-good-enough life better. Since, as I argued in Chapter 1, a condition of the good life is that it should have some content, it seems at least possible that improvements of a good-enough life could be provided by enriching the content thereof. But the field of areas of research that enrich our lives in some way – and so make it better – is very wide indeed; and this is not just because some research that looks to be otherwise useless may actually be indirectly useful (as would be true of the techniques to grow bone on a lattice that I explored in Chapter 5): it applies to things that genuinely are otherwise useless. So the way would appear to be open for people engaged in activities like planetary science or palaeontology to claim that what they are doing counts as being beneficent, *and for this claim to have some mileage*. A researcher might genuinely believe, with more than a small hope of

accuracy, that what she is doing will benefit either presently-existing, or future, people; and it's not at all inconceivable that she might be right. The enrichment of a life is an aim worth pursuing; if our concern is with the good life, then such a claim should not be all that hard to defend.

Crucially, once we have admitted that there might be benefits to be had in terms of enrichment and improvement of human lives even from research that we might dub 'sterile' – by which I mean that it will not lead directly to benefits beyond those that it brings itself – there is only a small step that needs to be taken to admit that the same benefits can be had from things that are not research at all, even though they, too, are sterile. (It might actually be harder *not* to make the step than to make it.) Living in a rich culture might be a benefit in its own terms. Claims that I made in Chapter 5 about the permissibility of strange body modifications serve, indirectly, to support this: if there are interventions that are defensible and worth doing notwithstanding their lack of 'objective' utility, then there might well be other things that are similarly defensible and worth doing despite being of limited obvious utility. Humanity as a whole is likely to benefit, albeit in an intangible sort of a way, by allowing people to pursue activities that are not directly beneficent, but which do enrich humanity in some way – culturally, for example. (This point echoes the idea mooted towards the end of Chapter 7 that the world as a whole might be better off by allowing individuals behave in a manner that is not obviously virtuous.)

The argument boils down to this: there are benefits to be had from sterile research in (say) physics; and if we accept this, then the same ought to be available in spades when it comes to research in philosophy, philology, or philately. We could stop doing all of these things without causing harm. But there might nonetheless be something valuable lost.[12]

And this point throws light on John Harris' claim that

> we can be simply curious if we like, but the justification for science and for philosophy is not simply in terms of the disinterested pursuit of knowledge but in terms of the benefits that such a pursuit can bring.[13]

If the justification for a research programme is in the benefits that it generates beyond itself, then it is plain that a sterile research programme, or artistic pursuit, will be at best less justified than one that is non-sterile. But what Harris minimizes here is the possibility – one that he nevertheless admits – that the disinterested pursuit of knowledge might be good in its own terms; its sterility may not be a problem. None of this makes any difference at all to the idea that it might be wrong to pursue the arts, or planetary science, or palaeontology, if doing so detracts from duties of rescue; and if that duty of rescue is best discharged by means of research (*pace* the argument in the next chapter that it almost certainly never is), then not researching will be blameable, and research *pro tanto* a duty. But since these conditions will be met in only a very few cases if any, it would probably be foolhardy to generalize too much.

This idea does prompt a couple of secondary questions. First, we might want clarification about whether putatively obligatory beneficence is properly directed at individuals, or at humanity in general (though coming down on one side or the other won't make a big difference to the argument here). Since there's a logical difference between humans and humanity, there might be things that benefit one without benefiting the other. (A tertiary question would concern whether, if it's properly directed at individuals, it's only currently existing people that matter, or whether – and to what extent – potential future people count as potential future beneficiaries.) Second, we might wonder whether the motivation of the people providing the benefit matters when assessing the discharge of a duty. By this, I mean that it is not obvious that a NASA scientist, or a palaeontologist, or an artist, would necessarily be motivated by beneficence to do what they do, even though it does create benefits; any benefits that do accrue would be a desirable side-effect. If we think that the duty of beneficence that we ostensibly face reduces to a duty to do things that could reasonably be expected to generate benefits, people like NASA scientists and artists could turn out to be morally in the clear, even if they are indifferent to outcomes that benefit others. If, though, intention matters, so that a benefactor has to create benefits for beneficence's sake, then the NASA scientist who is not motivated by benefit, but by (say) wonder, would still count as morally lacking. But, by this token, then so would the person whose discovery of a cure for HIV – who, that is, has done groundbreaking bioscientific and medical research – is motivated by the desire for a Nobel prize or a bare fascination with problems of biology. Meanwhile, there seems to be no reason why an artist couldn't claim to be working for the sake of beneficence.

## The argument from incommensurability

But maybe we're getting ahead of ourselves here. After all, in his articulation of the duty to research, Harris writes that

> there is a clear moral obligation to participate in *medical research or any research that has a reasonable prospect of enhancing our lives or our material condition in certain specific circumstances*. This obligation is importantly not confined to purely therapeutic research but also involves all beneficial research.[14]

The emphasis is added by me. Suspicious as we may be that this restriction on the nature of the obligatory research is arbitrary and imposed by *fiat* – especially allowing that the argument from Chapter 3 about enhancement not being a duty is sound – if we take it in good faith and on its own terms, it does seem to provide us with a way to rank the benefits of a given piece of research: the obligatory things are those that have a reasonable prospect of enhancing our lives or our material condition in certain specific circumstances. And, on

similar lines, when Harris says that the justification for research is not simply' in terms of the disinterested pursuit of knowledge, that does at least suggest that the disinterested pursuit of knowledge might be good for its own sake, irrespective of whether or not any real difference is made to the world or the lives of actual people. It so happens that, for Harris, the importance of such speculative endeavours is trumped by that of medical research; but that is not to deny it's good, any more than a gambler's King of Hearts being trumped by his opponent's King of Spades would allow us to deny that the King of Hearts was a king to begin with.

Yet even if such hierarchicalization has an intuitive attraction to it, this attraction does conceal a puzzle: by what standard is medical research more important than astrophysical or archaeological research? And by what standard is any research preferable to other ways of filling the day that are not research based at all?

It's worth restating here the working hypothesis that, if research is a duty, it is so because of the good that it will generate; but from that we ought to infer that there will be nothing that is capable of generating the same amount of good – the same improvement in people's lives. For if it were to turn out that some other, non-research, activity was capable of generating the same amount of good, and it's beneficence that is our driver, then that would mean that our duty was not to research, but to do that other thing. Maintaining the privileged position of scientific research requires the provision and use of some account by which we can say that what it provides is of such a value that there is no competing moral reason to do anything else – and it's hard to be sure that we could ever have such an account. It seems likely to me that the good that might be generated by further research, and the good that might be created otherwise, are qualitatively different. And if this is true, then the idea that we ought to choose one rather than the other on account of its ability to provide welfare is in serious difficulty.

Let's imagine someone who, at a certain stage in his life, is given the choice between becoming an artist and becoming a research scientist. Call him Ernest. Ernest is assured that the rewards that accrue to him will be about the same either way; so at least as far as he is concerned, he might as well flip a coin to decide which option to take. He realizes, though, that his choice will have an impact on the world around him; being a magnanimous sort, he decides to choose the option that could reasonably be expected to generate the best outcome. He now finds himself confronted by the question I asked a moment ago: By what standard should he assess the options? What's the metric by which the desirability of each possible future world is to be measured?

This is not an empty question. It makes sense to talk about two possible treatments for an illness in terms of which is better; that can be measured in terms of things like QALYs. It might even be possible to apply something like

this rubric to different areas of research. For example: if we're ever to have commercially viable nuclear fusion reactors, this will probably require a great deal of work from materials scientists to identify what materials we should use to make the things.[15] We could, in principle, calculate the QALYs generated by fusion research, and compare them against the QALYs generated by malaria research, and use that to decide how we should spend our money. (Naturally, if a QALY score isn't to your taste, some other metric might serve. The point stands that, if you're going to say that one ought to $\phi$ rather than $\psi$, you're going to need some way of establishing that $\phi$-ing is better.)

But not all research will fit into this rubric. The calculation of the Hubble Constant will never save any lives – but since that's not what draws people to understanding the universe in the first place, this seems like the wrong stick with which to beat the enterprise, if beat it we must. And just as there're goods to be had – in terms of enriching human lives – from non-biomedical or bioscientific research, there seems to be scope for similar goods to be had from activities that really aren't research at all. Nor does it bespeak a failure of rationality to prefer something useless over new bioscientific interventions. And just as a low QALY score isn't enough to demonstrate that astrophysics isn't good, so it won't be enough in respect of other things.

For example: imagine that Silas is a wealthy man who contracts a deadly but slow-to-progress illness like multiple sclerosis. Silas believes that he could fund research into a treatment, and that he would perhaps benefit from it himself, but decides that he will nevertheless sink his resources into the construction of a grand mausoleum for himself and family. This is not necessarily an irrational preference so much as a non-rational one: he simply has a vision of the kind of world he would want to create – a world informed by the idea that mausolea are important, and cures for MS are not so important. On the other hand, it is not hard to see why people might think his preference somehow indefensible, and seek to persuade him to do something with his cash that *would* help MS research or something of the like. And this would certainly be so if we're going to run with the twin assumptions that beneficence is a duty and that supporting research can be a way to discharge that duty – assumptions with which I'm happy to play along for the sake of the argument. Building a mausoleum is not obviously beneficent, and if building it comes at the expense of discharging that duty, then it might be positively wrong.

But now imagine that, instead of being a miser dedicating his funds to memorial architecture, Silas decides that he's going to spend it on endowing a fund that would pay for theatre trips for disadvantaged kids from the other side of town, or a piece of public art, or something like that. Here, the idea that there's a duty of beneficence doesn't seem nearly so powerful when it comes to constructing an argument about his obligation to support research instead, because Silas could claim to be acting beneficently: he thinks that humanity

*sensu lato*, or some disadvantaged slice of humanity, will benefit – and that benefit was what motivated his bequest. Had Silas sufficient belief in his own talents, he might be tempted to say the same thing in relation to his own pursuit of art rather than science. What would have to be shown however the matter is viewed is that he is not acting in the most beneficent way possible – but to establish the importance of that would require showing not only that there is a duty of beneficence, but a duty of *maximal* beneficence, and also that there is a clear manner of comparing the benefits of each option. It's this latter criterion that represents a particular weakness for the claim that there's a duty to research.

It's a weakness that won't be strengthened by having it pointed out that a new treatment for MS is probably important to other people. After all, the fact that something is important to one person doesn't mean that it is important objectively, or that others ought to adopt the subjective sense of importance. (That Stoke City FC should win the Champions' League is probably important to some people; it doesn't mean that it's important, or that anyone else has a responsibility to realize the dream.) Moreover, expanding the cultural horizons of the least well off is important to Silas; and it would appear that it's at least as important to him as MS research is to others. He doesn't have to deny that researching new MS treatments is a good thing; he might even admit that he's idiosyncratic in his preferences. It's just that research doesn't exert a particularly strong pull on *him*. Besides: if others' desire for a cure ought to influence his spending decisions, it's not easy to see why his preferences oughtn't to sway theirs as well, giving him the same complaint against them that they have against him. Merely pointing out that something is valuable to someone isn't normally enough to make a strong claim about others' obligations in respect of that something.

(Along similar lines, the anti-hero of BS Johnson's novel *Christie Malry's Own Double-Entry* contracts cancer, and berates the book's author for his choices:

> "In any case," he said, almost to himself, not looking at me, "you shouldn't be bloody writing novels about it, you should be out there bloody doing something about it."[16]

(But Malry – it has to be said – is hardly a disinterested observer. The fact that he – or anyone else – would prefer others to be doing something about his illness doesn't necessarily carry all that much weight when it comes to deciding what they should do, and his preference does little to refute the idea that there might be a qualitative difference between the goods provided by different sorts of activity. A person can think any number of things worthwhile within the context of his own life, and that some of them will benefit humankind to boot.)

The point is this: even if we think that there's a duty of beneficence, and are satisfied that this duty when unpacked amounts to being a duty of

maximal beneficence, working out what is maximally beneficent may not just be hard: it might be impossible, because the goods generated by $\phi$-ing may be qualitatively different from those generated by $\psi$-ing. If the benefits biomedical or bioscientific research are qualitatively different from the benefits of planetary science, ichthyology, or poetry, then it won't be possible to tell the poet that she ought to have done otherwise simply by pointing out the benefits that could have been accrued. If we happen to think that bioscientific research is the most important thing we can do, then this is fair enough: but it looks like an aesthetic, rather than a moral, claim that we're making.

And lest it be thought that I'm making too bold a claim here about the incommensurability of goods, I can afford to weaken it, and simply say that the onus would be on the defender of the putative duty to research to show that there is some metric by which they could be compared. Whatever metric is suggested will have to be carefully chosen, though, to ensure that it packs the normative punch that is required: it's not going to be enough simply to show that, as a matter of fact, more people do happen to value research highly than value non-research highly. This is because what people *do* prefer and what they *ought* to prefer are two different things, and it's what they ought to prefer – between some research activity and some non-research activity – that's at the heart of what matters here; and if we could show that, there'd be no obvious need to wonder what people do actually prefer.

There's a couple of other points of merit in this context. The first is that establishing a metric by which we could compare the importance of biomedical or bioscientific research and other activities would potentially create one great hostage to fortune, since the defender of the duty to research would still have to be sure that a given course of action – in this case, research – genuinely does provide more benefit than anything else; if it doesn't, then to privilege it amounts to an endorsement of the less good over the better. Moreover, if the privilege – and the concomitant duty claim – is to be robust, it ought to apply not just to affairs as they are now, but to a range of states of affairs – and all of them, if at all possible. If it should turn out that there is a possible world – including ours in the future, depending on the marginal utility curve of further scientific research – in which the benefit provided by non-scientific research is greater than that provided by scientific research, there will be no obligation to pursue science, and a potential obligation not to. Put another way: scientific research's position would appear to be only contingently privileged. This contingent privilege is precisely what one would expect, of course, if one adverts to the claim that the duty to research is just one version of a more generalized duty of beneficence. But to make this kind of move would be suicidal for anyone who wants us to take seriously the idea that there is a general duty to research rather than a set of duties to pursue certain particular areas of research, each of which would require separate argumentative support.

The other point is this: whatever metric we might generate to compare goods would – it appears – require a good deal of philosophical argument and analysis. This is as much as to say that establishing a duty to support scientific research is parasitic on endorsing something other than scientific research as well. Being able to say that we have a duty to pursue the good through science requires a conception of the good the like of which science is incapable of providing.

## The argument from anthropology

Scepticism about the means by which we compare the goods accrued from a range of activities is supported by an appeal to a kind of anthropological relativism.

The appeal emerges when we note that there exist several cultures around the world that are minimally technological, and appear not to have conducted any significant scientific research for thousands of years: for example, there are tribes in the Amazon that have little technology beyond the bow and blowpipe. In some cases, technological development appears not only to have halted, but to have regressed; Timothy Taylor talks about how the stone tools used by aboriginal Tasmanians at the time of European exploration were demonstrably as simple as those used by chimpanzees[17] – but, more importantly, less sophisticated than those Tasmanians' own ancestors': by the time the Europeans arrived, they went naked despite the climate being not much warmer than that of the colonists, and some tribes seem to have abandoned permanent shelters altogether in favour of basic wind-breaks. They had – so it looked – even lost the ability to make fire for themselves, relying instead on fire-sticks that were kept alight from previous fires and carried from place to place.[18] To the eyes of the Europeans who met them, they had done nothing very much to advance themselves beyond a state of nature. Their world was more or less completely unimproved.

On the face of it, then, and assuming that there is a responsibility to pursue the good (perhaps through the pursuit of biomedical or bioscientific research), we have to ask how it could be that certain whole cultures are so fallen as not to discharge this duty, or even to abandon advances already made. Echoing a question asked by Harris and Holm,[19] must we assume moral turpitude in non-technological cultures?

This question invites a range of possible answers. One is that these cultures are simply morally lacking. A duty such as the duty to rescue drowning children when doing so is easy is one that we think applies universally: anyone who failed to attempt it without a very good reason indeed would be blameable; and the putative duty to research insofar as it's a means of rescuing people from a harsh life could easily be cut from the same cloth, and supporting

or participating in research or some sort of technological advance could be a requirement of anyone, anywhere. Yet there's clearly something a little queasy-making about writing off whole cultures as morally lacking, at least as a first resort. For while the practices of the past and of different cultures may strike us as morally indefensible now – we no longer burn witches, and we are confident that this is a moral improvement – the idea that a culture could make this kind of basic moral mistake for so many thousands of years brings us uncomfortably close to the idea that the people living within it were somehow constitutionally incapable of discharging the duties incumbent on the rest of humanity.

A second possible answer is that non-technological cultures are marked by a succession of attempts at technological improvement, pretty much every single one of which has failed. This avoids the rock of suggesting that some groups are constitutionally morally lacking, but lands us squarely in the hard place of suggesting that some groups are constitutionally good but stupid. Besides: given the apparently obvious advantages that even a small technological advance would give a competing culture, it would remain a mystery how it is that these recalcitrant cultures survived as long as they did, or – even if they never developed anything on their own behalf – why they didn't simply steal technological ideas of one sort or another from their neighbours.

A more interesting possible explanation is that they actually have improved their lives and the world around them in some way: it's just that their standard of improvement does not take them down a technological route. And this is the hypothesis that Taylor articulates in respect of Tasmanians. If, he suggests,

> you can adapt psychologically and physiologically to a life without clothes in a high latitude, then you are freed from an awful lot of sitting around, weaving, sewing, mending, and, of course, getting damp and cold.
>
> No bone tools means no awls means no clothes means no pockets means nowhere to keep tinder and fire-making kit; that means no fire-making, which means carrying fire all the time. That means quick fires whenever you want, which means it is okay to have no clothes. Carrying fire all the time means it is safer not to wear clothes, so you don't catch fire by accident. It also reduces what can be carried. No composite tools means no axes, so no log boats or all-weather craft, which means no fish – but with lots of sudden storms, why risk orphaning the children when seals can fish and you can eat seals. Inshore, you can grease up and dive for lovely stuff. Naked, of course, because you don't want damp clothes. You want to get dry and warm as fast as possible by the fire, and eat abalone, wallaby, and tree-harvested possum.
>
> Without a complex toolkit to lose, or surpluses to be stolen, or clothes to dry, you don't really need a house, and since you carry fire everywhere, the risks of burning one down would be high. With no houses and little personal property, there's not much hierarchy. Accounting is unnecessary, so you don't need writing or numbers. You have no maps, but you are not lost. You know where absolutely everything you need is. And because you don't have to look after it, you can get it when you want it.[20]

Tasmanian technology was honed to the bare essentials, to be sure – but this did not indicate a failure, so much as a certain rugged efficiency. For sure, they may have suffered hardships that other groups did not suffer – but, on the other hand, others suffered hardships that they did not. Importantly, and only tacit in Taylor's account, it may be that the hardships facing the Tasmanians were only particularly stark to the Europeans; and, equally, it's possible that they'd have considered apparently unobtrusive aspects of the European lifestyle – the requirement to labour, to maintain housing and clothing, and the whole mechanism of capitalist accumulation – as a ridiculous and needless burden. Their way of life may have been Spartan – and it does not (at least to my eyes) look particularly attractive; but it is a heck of a leap from saying that to saying that by not developing certain technologies, and abandoning others such as boat-building and fire-making, there was any moral failure.

Rather, what they had was a – fairly radically – different way of life. Presumably, it was not one in which people felt any particular need for a great deal of technology – or, implicitly, a great deal of medical innovation. There's no real need to wonder about a failure to discharge a duty of beneficence if the kind of life that is currently lived is accepted as being good, or at least good enough. Indeed, if Taylor is correct in his assessment of the possible advantages of becoming less 'entailed', we could even mount a case for thinking that a culture that divests itself of certain notional badges of development does so in response to a nudge from a beneficent invisible hand.

The same applies for other cultures that have not, historically, followed the same developmental route as the urbanized and technocentric West. We can argue about whether the kind of life adopted by a culture (and its members) really is the best one; but those arguments are not massively important here. What matters here is the claim that a putative duty of research is best seen as simply an aspect of a much wider duty of beneficence. Research is one of the ways – and only one of them – in which the culture with which we're most familiar has chosen to discharge its duty of beneficence; but 'beneficence' makes sense only in the context of a claim or set of claims about the kind of life worth pursuing. With a different account of the good life, and of the ultimate justification for any action, a different measure of what would count as beneficent might dictate non-research pursuits. So while it is in a sense possible that the failure to pursue biomedical or bioscientific research is attributable to some colossal collective moral failure by the members of non-technological cultures, that is not the only explanation, and it isn't really a very good one.

## Ecology and economy

Though it's a little bit of a distraction, I think it's worth spending a couple of paragraphs talking about another aspect of a putative duty to research – an aspect that covers what we might dub 'ecological' reasons to research.

In a nutshell, the argument goes something like this: the world is a finite place, and the 'economics of coming spaceship Earth' are such that we can be reasonably sure that current lifestyles are ecologically unsustainable and environmentally damaging; this means that there is a chance that the lives of future generations may be avoidably blighted. Bioscientific research may give us a way to avoid any such blight. For example, synthetic biology takes what we have learnt about genetics and offers the promise that we could design and build microorganisms that would be able to produce petrochemicals in a near carbon-neutral way, clean up chemical spills by digesting toxins, and so on. In other words, there's a range of ways in which bioscience might help us neutralize the risks generated by our current lifestyles.

Yet the tacit assumption here is that our current lifestyle is either inevitable or sacrosanct. It might be that research in the biosciences is required *if* we want to retain the kind of economic model to which we're used – but this imperative is only hypothetical, and there's nothing to say that we could not permissibly opt for a steady-state economy, or even one in which we forego certain aspects of modern culture and thereby replace one kind of plausible good life with another. I argued in Chapter 1, after all, that culture can play a significant role in determining what counts as a good life. It does not strike me as impossible in principle that some future culture might decide to opt for a different lifestyle – one that is not so polluting. Historians in this future world may look at our lifestyle and wonder how on Earth we ever found it tolerable, preferring as they do fewer gadgets but more free time. Or it might be simply that the putative obligation to carry out research to mitigate the impact that we have on the world is matched in its benefit by a putative obligation to have fewer children, thereby reducing the human population to sustainable levels. If we're considering obligatory changes to lifestyle, after all, that might be the kind of thing we should countenance.

I am not arguing here that there is anything morally troubling about the current economic model under which we live. (There might be, but it's not what I'm arguing.) What I am saying, though, is that there's a set of choices to be made that is often forgotten.

The point is this: the pursuit of things like synthetic petroleum is, arguably, desirable *subject only to a range of unexamined claims about the good life*. Insisting on more research to facilitate the continuation of the kind of life we currently lead is one thing; but it won't tell us anything about duties unless and until we've accepted the idea that continuing the kind of life we currently lead is a *sine qua non* of promoting flourishing. The dilemma that forces us to choose either embracing further research in order to protect the quality of future lives, or risking that those lives will be of a lower quality than we would find acceptable is illusory; it is in fact a trilemma at the very least because there are more options than we might think. Research in the biosciences may be *a* way out of foreseen problems; but (even allowing that there is an obligation to avoid these foreseen problems) unless it can be shown to be the only way, the chance that it'll turn out to be a duty on this basis looks fairly small.[21]

## Is there a duty to research?

Is there, in the end, a beneficence-based and wide-reaching duty to research? While it'd be foolhardy to insist that there certainly isn't, I think that the answer must be that there is at the very least no obvious reason to think that there is. There are certainly reasons to research, but they do not seem quite to tip over into the realm of obligations.

It is plausible to think that the reasons to pursue some kinds of research are more powerful than the reasons to pursue research of other kinds. Notably, research into techniques that would make the difference between a less-than-acceptable life and a good-enough life might strike us as being more important than research into techniques that would make the difference between a good-enough and a more-than-good-enough life. This distinction tracks the T/E distinction as outlined in Chapter 1. Still, being able to describe research as therapeutic or 'therapogenic' will not be enough on its own to legitimize an ascription of obligation.

One of the reasons for this is that, even if we think that therapy is an obligation – even, that is, if we think that we have a duty to intervene to make a person's life or some aspect of it at least good enough – research is not the only, or the best, way to go about it. I'll elaborate on this in the next chapter; but the thought is that though (say) research into malaria treatments might very well be therapeutically-directed, research is by its nature an uncertain enterprise. As such, it might not be the best way to provide relief right now; we might always have a better reason to do something else. Some of the entries on the list of other things that we could do will be related to the relief of symptoms, rather than the condition itself; sometimes, they will be more along the lines of simply attending to the sufferer in wholly non-medical ways: the word 'therapy', we shouldn't forget, has very many senses, not all of which denote cure.

So while research on a vaccine for malaria could be admirable, on the grounds that suffering from malaria is quite probably the kind of thing that reasonable people would hold to be antagonistic to the good life, because research is the work of years, eradication of the parasite that causes malaria might also fit the bill here as the move required to raise standards of flourishing to an acceptable level – and it might be a more effective way to get rid of the disease anyway. To give another example, it might be that research into creating viable eggs from other cell types, which would allow at least some infertile women to have genetically related children, would not be obligatory, on the grounds that having genetically related children is not clearly part of the good life – indeed, such research might actually make some people's lives *worse*, on the basis that it appears to confirm the fallacious and potentially harmful idea that a family is necessary for the good life, and genetic relationship with someone is necessary for them to be a part of your family. (Similarly, recall Brian from Chapter 1: his infertility is a problem in want of a solution only for as long as we treat

as a given that having children generated from gametes that originate in his own body is all that important. If this is denied – as it likely can be – then the imperative to 'cure' the problem is somewhat diminished.)

So where does this leave us? The thesis of this chapter has been that the most obvious reason to think that there might be a duty to research would be to tie it to a duty of beneficence (and so, indirectly at least, to an idea of what constitutes the good life). But it's not clear that a duty of beneficence really obtains.

More importantly, if it does obtain, that it's not at all clear that it's strong enough to underpin a duty to support more-than-minimal research; if it goes any further than that, then it seems to collapse into the kind of thing that could equally well support all manner of things that are either what I have called sterile research, or are not research at all. And if the duty of beneficence allows us to pursue sterile research or things that are not research at all, it would not obviously be blameable to spend our time other than bioscientific research; and if it's not blameable not to do that, then it's hard to see that it was a duty in the first place.

Still: this is not the only kind of argument that might be mounted to establish that there's a duty of research. In the next chapter, I'll turn my attention to a few other candidates.

## Notes

**1** Harris, 2005, repr. in Harris, 2007.
**2** Harris, 2007, p. 186.
**3** Stjernschantz Forsberg et al. 2013, *passim*; cf. Brassington, 2013a, *passim*.
**4** Mill, 1998, p. 172.
**5** Ibid., p. 169.
**6** Kant, 1993, 4:423.
**7** Kant, 1996, 6:453.
**8** Williams, 1973, pp. 116–17.
**9** Kant, 1993, 4:430.
**10** Jonas, 1969, p. 230.
**11** Zhang et al. 2010, p. 1077.
**12** In the course of a conversation with Sarah Chan, I was reminded of the metaphysicians of Tlön in Borges' story, who 'are not looking for truth, nor even for an approximation of it; they are after a kind of amazement' ('Tlön, Uqbar, OrbisTertius', in Borges 1998, p. 24). Who would begrudge them that?
**13** Harris, 2007, p. 186.
**14** Ibid., p. 192; slightly modified.
**15** As I understand it, we are at best uncertain about how to build a reactor the fabric of which could stand up to the neutron bombardment it would undergo.

**16** Johnson, 2001, p. 180.
**17** Taylor, 2010, p. 33.
**18** Ibid., *passim*.
**19** Harris and Holm, 2003.
**20** Taylor, 2010, pp. 50–2; slightly modified.
**21** One thing that I don't have time to consider here is the Parfittian possibility that there is no obligation to avert a sub-optimal future in which lives will be unnecessarily hard, but still worth living. If it's possible that the people of that future could not have existed in some alternative future, it might be that we don't wrong them by not taking more care of the environment.

# 9

# Bioscience and the Duty to Research, Part 2: Non-Beneficent Arguments

In the last chapter, I expressed doubt about whether there is a duty to research for the sake of a beneficent desire to improve the lives of others. Demonstrating that such a duty obtains would mean demonstrating that beneficence is a duty, and this is harder than it may seem. But even if beneficence can be shown to be a duty, we can't be sure that it's a duty of the right sort – that there is only a limited number of ways to discharge it. Rather, what duties of beneficence there might be seem to be able to accommodate a fairly broad range of ways to discharge them. There's no particular guarantee that research would figure among them.

The point of this chapter is to look at what I take to be the strongest of the remaining candidate arguments for there being a duty to research. As it happens, all have been articulated by John Harris in the past few years, and so I'll take his arguments as my touchstone; but the counterarguments I'll be presenting ought to be recognized as having an eye on a wider horizon.

## Formulating the duty to research

A paper published by John Harris and Søren Holm in 2003 about the participation of children in research takes as a central premise the idea that there is a duty to participate in, or at least to support, medical research. The point of that paper was not to argue for the premise, but a couple of lines do set the scene for what might be developed as arguments substantiating it:

> If a given research project is well designed and likely to lead to important knowledge that will help persons in the future, can there be a moral obligation to accept being a research subject if asked (or to volunteer)?
>
> The answer to this question is an unequivocal "Yes". Two separate lines of argument lead to this conclusion. First, it follows from our basic moral obligation to help other people in need (or our duty of beneficence). . . . Second, the obligation also flows from a standard Rawlsian "free-rider" argument.[1]

Within the context of the paper, Harris and Holm don't say much more to bulk out these ideas; and I hope that I did enough in the last chapter to show that the beneficence claim is not compelling when interpreted in a simple manner. But what about the other aspects? How might appeals to free-riding, for example,

generate claims about a duty to research? Obviously, more must be said if the kind of thing towards which Harris and Holm seem to be nodding is to be taken as compelling.

The arguments are filled out to at least an extent in a 2005 paper in which Harris claims that there is an obligation to support biomedical research, perhaps by participating in it. (For the sake of space, I'll refer to this as a duty to research, or a duty to support research, as the context demands: the same general thing should be understood by all these formulations). He does so by appealing to what I think are three quite separate arguments. Two of these are diachronic, drawing a link between research done in the present, and past or future benefits; the other is synchronic, concentrating on the 'present tense'. The synchronic argument, which I shall call the *Free Rider Argument* (FRA) rests on the intuition that it behoves us not to take for granted the fruit of others' labours. The argument is superficially similar to the first diachronic argument, which looks back to the benefits that we have received from society so far to ground a claim that we owe certain things back, and can discharge the obligation by supporting research; I'll call it the *Argument from Filial Piety* (AFP) here. Such is the proximity of the FRA and the AFP that they might be held to be versions of the same argument; but I think it's worth keeping them separate, even if the separation is slightly forced, for the sake of the counterarguments I want to present. The other diachronic argument looks forward to ground a claim that we have an obligation of beneficence to future generations that mandates support for further research; I'll call this the *Argument from Rescue* (AR).

Though the AR is clearly related to beneficence arguments, I think that it merits separate consideration, because duties of beneficence seem to me to mandate a particular attitude towards others' wellbeing in agents, while the focus of a duty of rescue appears to be directed towards specific cases, and doesn't really require that we spend too much time thinking about the welfare benefits to be generated: that is, we might make someone better off by rescuing them, but that is not *why* we rescue them. Besides: there are arguments against the use of research as a method of rescue that hold independently of whether or not the beneficence claim works.

Finally, in Chan and Harris' 2009 restatement of the duty to research, a hint is dropped that there is a fourth argument that might be produced to establish a duty to research; I'll call this the *Prevention and Causation Argument* (PCA). It's with the PCA that I'll begin my analysis, before returning to the 2009 paper towards the end of the chapter.

## The prevention and causation argument

The PCA is quickly articulated, and quickly disposed of. The nub of it is captured in Harris' claim that 'a decision to withhold benefit is always damaging', [2] and

echoed in Chan and Harris' claim that '[f]ailing to prevent harm is as effective a way of ensuring that harm occurs, and hence as morally reprehensible, as doing harm directly'.[3] Research, the argument goes, is harm-preventing, and so a failure to research is a failure to prevent harm. And if it's true that a failure to prevent harm is an effective and reprehensible way of ensuring that it occurs, it must follow that a failure to research is comparably reprehensible. This doesn't quite amount to a claim that research is a duty; but it does come pretty close. However, it should be clear that if it happens not to be true to say that withholding benefit is damaging, or that failing to prevent harm is as effective a way of ensuring that harm occurs as doing harm directly, the argument will collapse. And it so happens that it is *not* true to say either of those things, and the argument therefore *does* collapse.

A couple of quick examples ought to show why. If it were true that failure to benefit were damaging, it would be true that I am making someone worse off by not making them better off. This is quite simply false. For sure, they are at the end of the day worse off than they might be in some possible world; but that is not the same as being *made* worse off. Not giving money to a beggar is not an attenuated form of stealing from him.

Nor is it true that failing to prevent harm is the same as – or even in the same moral league as – causing it. An anecdote will show why. A few years ago, I had the good fortune to visit Victoria Falls. There is a path along the cliff, and it is – or was at the time – possible to walk right to the edge, there being no barrier to prevent hapless people who got too close to the spray-wet edge from slipping and falling to the rocks and crocodiles far (horribly far!) below. Doubtless, such a fall would be harmful. It would be absurd, though, to say that the custodians' failure to put railings in place *ensures* that that harm occurs: my writing this now proves as much. Ensuring that people came to harm would have required more than not fitting a barrier – it would have needed someone to push them, or to jump out from behind the bushes shouting 'Boo!', both of which seem to be much more serious matters than not having put up railings.

This is not to deny that there might sometimes be something morally problematic about leaving the world as it is; we may have an obligation to give to the beggar on the street. But this is not because we harm him by not doing so; and failing to improve his lot, or to improve the world in some other way, is not qualitatively the same as positively making an alteration for the worse. Moreover, while there may be a reason to act in such a way as to prevent harm, there may at the same time be other reasons not to. To keep with the waterfall example: there are plausible aesthetic reasons not to fit a barrier. If these aesthetic reasons are sufficiently strong, they may suffice to counterbalance the reason to fit one; and, this being the case, there would be, again, no such duty. Of course, if one thinks that prevention of harm is the sort of reason that always comes up trumps, then one is free to argue for the claim – but it does need an argument. Still: the important point that failing to prevent harm is not a way of ensuring that harm occurs would stand; and while it may

be reprehensible in some circumstances not to prevent it, it does not follow that it is *as* reprehensible, or reprehensible for the same reasons and by virtue of the same argument. If there is a duty to research, it is not held together by the PCA.

## The argument from rescue

Though a failure to prevent harm is not the same as causing harm, preventing harm might conceivably be presented as a form of rescue, and therefore morally required. Harris claims that '[w]here our actions will, or may probably prevent serious harm then if we can reasonably . . . we clearly should act'[4]; elsewhere, he points out that '[i]t is surely unethical to stand by and watch three million people die this year of Aids alone and avoid taking steps to prevent this level of loss'[5]. And since medical research provides a way in which we can prevent serious harm and suffering, it is therefore unethical not to carry out research. Just as in Peter Singer's famous example it is inexcusable not at least to attempt to rescue a child drowning in a pond,[6] so it is inexcusable not at least to attempt to ameliorate the suffering or prevent the premature deaths of others when this comes at a proportionately low cost such as a few pounds given to a research charity or a few hours taking part in an experiment. If we have a duty to rescue the drowning child – and we do – we have a concomitant duty to rescue people in other ways.

The argument is clearly stronger than the PCA, not least because it gets to the point – the one about avoiding and militating against harm – more directly. Expressed slightly more formally, the AR looks a bit like this:

(1) We have a moral duty to effect a rescue where possible
(2) Saving lives is a form of rescue
(3) Research saves, or can save, lives
(4) Therefore research is a form of rescue
(5) Therefore research is demanded by the duty of rescue

All the same, though, the argument is not compelling. This is for a couple of reasons: first, the truth of (4) is open to doubt; second, even if (4) is true, the likelihood that it is the only or the best way to save a life seems to be low. This would mean that we probably have a moral reason to prefer other things ahead of research; but since a duty is that action the omission of which is blameable,[7] and it might be preferable on a given occasion to do something that isn't research, then research cannot be a duty; hence (5) falls even if (1)–(4) stand.

It doesn't take all that much to show that (4) is unconvincing. The first step is to clarify what it means to effect a rescue. The problem with (4) is that accepting it will require a radical departure from the everyday understanding of

the word 'rescue'; and yet – I'll claim – the only reason to make this departure is to keep the argument going.

In Chapters 3 and 4, I drew a distinction between saving a life – that is, rescuing someone – and avoiding death by means of life extension. In everyday language, to rescue someone means to save them from some present danger. Thus I rescue someone from drowning if, and only if, they are about to drown and I intervene to prevent that outcome. As per the argument outlined a few chapters ago, I do not rescue them by closing the swimming-pool prior to their arrival, or by teaching them to swim. Things like that might reduce the probability of harm or danger, but taking steps to ensure that a person does not face a danger is logically different from saving them from it. You cannot rescue someone who is not in difficulty; the most you can do is to prevent their getting into difficulty to begin with. And it is unsustainable to insist on moral identity in the face of logical difference.

Research, I submit, is more like the prevention of difficulty than its remedy. To see why, let's remain next to the pool. If we witness someone in difficulty, we may have a *ceteris paribus* obligation to wade in. But one would have a hard time making a case that research could represent any kind of parallel: it simply does not save people from imminent danger. Even with due deference to the fact that people who have got into difficulties in the water may benefit greatly from more effective methods of artificial respiration, the duty to rescue would not be discharged by undertaking research into those methods when confronted with someone who actually is drowning. And it is correspondingly hard to see how *any* research would satisfy the moral demands that the duty to rescue articulates.

There's another reason to reject (4), which is much more practical, and which harks back to the Endless List Objection that I raised to claims about enhancement in Chapter 3. The thought is this: that when I rescue a child from a pool, it stays rescued. There is an agreed standard by which we measure success, and success is measurable and achievable. The same applies, with appropriate modification, in respect of other kinds of rescue. If there is a famine, we supply grain, and measure success in terms of the number of people whose malnutrition is ameliorated or – at worst – stabilized: those who do not become more malnourished we can claim, with some justice, to have rescued. We might have to launch another rescue very soon; but the point is that we have a criterion by which we can tell that our duty to rescue has been discharged. This is not so clearly the case in respect of research. Because we would not be saving anyone from imminent danger, there is no obvious way to tell when our duty to research has been discharged – and so no way to tell when we had done what was morally required. Perhaps the thought is that there will always be an imperative to do more research – but this is implausible. As I hinted above, and as we'll see in more detail in a little while, one of the reasons for the implausibility is that there will be several situations in which

research is ruled out by moral demands: it would not only be the sub-optimal thing, but the *wrong* thing to do.

A last-ditch attempt to save (4) might be mounted by insisting that all of this is just semantics: maybe people who think that a duty to research can be established by means of an appeal to a duty to rescue simply have in mind an unconventional meaning for the word 'rescue'. This supposition might appear to make good the deficit in the argument. However, the problem with taking this approach is twofold. First, it redefines the word 'rescue' by stipulation; and since there does not seem to be any independent reason to think that such redefinition is warranted – what warrant there is comes wholly from the need for the word to carry a certain load in the argument – it looks to be rather *ad hoc*. I think that this consideration alone is enough to derail any appeal to the AR. Still: if we can perform the linguistic contortions without wincing too much, and are disposed to accept that 'rescue' can mean 'neutralise a danger in advance', (4) might – at least for the sake of the argument – live to fight another day.

But the second problem is that we might still legitimately think that rescue-by-research is a good deal less urgent than rescue in the more conventional sense, just because of the facts about who is in imminent danger of perishing. That being the case, a putative duty to research may well never be anything more than putative, and (5) would come in for attack.

One of the important points that underpins this attack on (5) is that there is only a certain amount of time, effort and money that a person could reasonably be expected to put into rescue. (Chan and Harris admit that a fair contribution to a project does not have to suffice to realize that project's ends, which suggests – assuming, plausibly, that there is no obligation to contribute more than is fair – at least the possibility that one may have discharged one's duty long before any material difference is made to the world.[8] Quite where one would draw this line is open to debate; but this doesn't mean that there isn't one to be drawn somewhere.) Now: imagine that 'rescue' can be quantified, and that each of us has five units of rescue that we can be reasonably be expected to distribute. This is a credit we can spend either on rescue by direct assistance (by which I mean things like providing grain to the starving, rowing the lifeboat, holding the ladders for the fireman, and so on), or rescue by research. It seems obvious that there would be nothing blameable about spending all of the credit on rescue by direct assistance. But if there is nothing blameable about not spending this credit on rescue by research, it is no longer possible to think of it as a duty, since (as I've claimed) a reasonable definition of a duty would seem to be 'that action the omission of which is blameable'. It would appear, then, that supporting research as a means of rescue is a duty at the most when there is no better way to improve the world. But what is important about this is that there will *invariably* be better ways to improve the world.[9] And this puts quite a big dent in any claim that research might be a duty when that claim is based on an appeal to rescue.

For one thing, research is by its nature an uncertain business. Let's assume that there is some standard by which we can measure the value of the contribution made by different kinds of rescue, and so compare them. I might therefore find myself in a position in which I could give £5 to give to Oxfam, or contribute £5-worth of research and so make the world a better place to exactly the same extent. The problem here is that rescue-by-research would always have to come at a discount because mere research does not guarantee the amelioration of anyone's life; indeed, even the very best science – methodologically sound science that takes full account of our best current understanding of the world – might be stymied by some hitherto undiscovered phenomenon. As such, it might turn out to be useless at achieving what it set out to achieve. And while it's true that the precise outcomes of any action whatsoever are strictly speaking uncertain, the problem is much more acute in respect of science than in respect of – say – fishing children from ponds or distributing grain to the starving.

What this means in practice is that responding to calls to research as a means of rescue will invariably be a very low priority; in fact, unless – *mirabile dictu* – we happen to live in a world in which there is no more immediate call on our moral attention, it will be something that we will always have a good reason to defer in favour of something else. It is trivially true that, if I am faced with a range of ways to rescue a drowning child, I should obviously opt for the one that is most likely to work; similarly, it is trivially true that, given a choice between saving one person from drowning or another from hypothermia, I should try to save that one whose survival I would be more likely to ensure. Presented with the choice between taking one person out of immediate danger, and contributing towards the *possible* rescue of another thanks to my participation in research on, say, a malaria cure, there is no contest. It is unfortunate if the person with malaria is doomed without the research (though that's unlikely: there are plenty of people in the developed world who enjoy decent lives even with the infection, because they have access to good medical attention even of it does fall short of a cure); but the AR means that, perhaps with a heavy heart, I must turn my attentions elsewhere. That which I ought to give to save others, I ought not to invest in research. Even if the development of the new drug is in its final stages, it is *still* overwhelmingly likely to be a more effective use of my time and effort to effect a rescue. This is why an appeal to rescue suggests that we ought to treat research as a very low priority. Of course, as soon as there is an effective drug that I could help supply, it might be a matter of indifference whether I opt to give that or grain. But in this case, of course, we are no longer talking about duties to *research*, because in this case the research has been done.

A possible retort here could be that, by providing a £5 sack of grain, I might be able to rescue a family of four. By contrast, even though my contribution of £5-worth of research is unlikely to save anyone in its own terms, it is, like the mouse's contribution to digging out the enormous turnip in the children's

story, a part of a much bigger effort, and thus is liable to be magnified many times. My contribution punches above its weight, for the development of a cure for malaria would rescue millions of future people. Unfortunately, there is no guarantee that the retort is going to do enough. The millions of future rescues that we might effect are millions of *potential* rescues, and, while they are indubitably desirable, there is no self-evident reason why I should not count the *actual* rescues that we could effect by paying attention to real emergencies happening right now as far more urgent. Moreover, while we might have an obligation not to *harm* the future, we do not necessarily have a duty to *benefit* the future: future malaria sufferers are not, at least as far as I can see, *wronged* by us when we spend our money on grain for the hungry.

The overarching point is that research is not obviously an efficient, or even effective, way of setting about the task of improving the world or others' lives. Things might conceivably be different if we lived in a world of unlimited resources, since we would no longer have to worry about the efficiency of our choices in terms of lives saved per pound spent. But resources are limited, and what resources there are could be used in ways other than research – we would not thereby be doing nothing. If one is won over by the idea that there is a duty to rescue, embarking upon or participating in a programme of research is not the best way to set about matters; one ought to consider the ways that are better. And if one thinks that there is a duty to research, appeals to rescue are not what sustains it.

## The argument from filial piety

The next kind of argument suggested by Harris to establish a duty to research that I want to examine is articulated in the claim that

> [w]e all benefit from living in a society, and, indeed, a world in which scientific research is carried out and which utilises the benefits of past research. It is both of benefit to the patients and research subjects and in their interests to be in a society which pursues and actively accepts the benefits of research and where research and its fruits are given a high priority.[10]

The claim is simple enough on the face of it: we have benefited from research, and should generate further benefit by participating more. Note that the strongest form of the appeal here is not that we should be aiming to generate benefit *per se* – that would fall foul of the objections to appeals to beneficence – but that we should be aiming to balance out the benefits that we have received and do receive thanks to others' research. The question that presents itself here is whether there really is any normative significance is of the fact that we benefit – that our lives are better or likely to be more flourishing than they counterfactually could be – from past research.

Establishing normative significance is no easy task. One of the problems that would face anyone who hoped to make the argument would be the need to explain how a description of my present condition translates into a moral obligation, or what the nature of that obligation would be. In Harris' hands, the argument would seem to run along the line that, because we owe the conditions of our present existence to scientific research, we ought at least to support ongoing research. However, it is fairly easy to construct analogous arguments to expose the problems of this line. For example, the battle of Edge Hill might have fulfilled one of the conditions necessary for my current existence, explaining how two distant ancestors came to be in the same place and conceive a child as a result. But this does not mean that I would have to endorse battles. A presently existing person does not face obligations *just because* he would not be where he is today without certain past events. Even if we could establish that I would not be here without some particular piece of scientific research, or research more generally, it would not follow that any obligation is generated in relation to further research.

The importance of this sort of consideration is not confined to arguments about existence: it is important in light of the weaker (and more attractive) claim that I might have obligations because of an *improvement* in my actual life, compared to that which I would otherwise have lived, afforded by technology; and it is important for the same reasons. There is any number of factors without which it would be plausible for me to think I would live a less good life. But it doesn't follow from the fact that something has contributed to the goodness of my life that I have any obligations. For example, one of the things that might reasonably be held to contribute to a person's flourishing is a reasonably wide circle of reasonably interesting friends. But whatever obligations one owes to one's friends, it wouldn't be a *quid pro quo* in return for that aspect of the goodness of one's life.

Moreover, so many factors have influenced and improved my – and anyone else's – life that, even if they *did* magically beget obligations, it would be impossible to meet all of them; it might even be impossible to decide which factors were the most important and therefore most worth trying to meet. Medical technology and research may well have contributed to the person that I am now and give rise to obligations; but this contribution will only be one tile in the mosaic, and I might have other obligations to meet before I get to the medical research facility. For example, I might decide, and (who knows?) decide *correctly*, that learning Latin at school had more of an impact on my current level of flourishing than had medical research, and that I thus have a more pressing obligation to spread its benefits. After all, I have never been seriously ill, and have no reason to suppose that I was ever in much danger of becoming so; but without having had to translate a passage about Quintus going to the forum to buy a copy of *The Republic*, it's possible that I would never have become interested in philosophy, and so not have ended up doing

what I do for a living now – something that, compared with some of the other possibilities (and knowing first-hand what it's like to work on a pot-bank), is a pretty good way to earn a crust. So, even if the AFP can generate obligations, there is no guarantee that they would be the obligations that the proponent of a duty to research wants: my primary obligation might be to endorse a classical education, since that's done the most good for me.

The most that the AFP could give is the claim that, because I would not be here or because my life would be measurably worse without certain things, I owe some sort of debt of gratitude to the past that made me the person that I am. But even here there is trouble, because we're violating the *is/ought* distinction. A bare description of all the factors that make me what I am is as poor a way of generating an obligation of gratitude to the past as it is of generating an obligation of research in the future without an additional rule outlining the proper attitude to one's past. But if there were such a rule for that, it would be unclear why there mightn't also be a straightforward rule or intuition describing a duty to research, making the AFP superfluous.

## The free rider argument

The final line of argument articulated by Harris in the 2005 paper – and, I think, the most sophisticated and powerful one – draws on the free rider problem. The free rider problem describes a situation in which one person benefits from the efforts of another to provide some good that both desire without making any serious contribution himself. The outcome is that, since everyone knows that he will be able to reap the benefit without carrying the burden, noone will be motivated to do the work, and the desirable thing is less likely to happen at all. To this extent, it's an economic, rather than a moral, problem. However, it is not hard to see that a moral dimension could be added, because the possibility of free riding gives each agent an incentive to minimize his contribution to some collective good; such minimization ensures that you'll reap all the benefits of some activity undertaken by others without making any sacrifices, and this looks to be, at least on the face of things, unjust.

The argument translates to research. The possibility that Smith may benefit from Jones' research, and *vice versa*, means that each has an incentive to wait for the other to do the work; this makes it less likely that the work will be done at all. And if either of them eventually does do it, the other will get the benefit of that research without having contributed – which seems to be unjust. It's the justice part of the argument, rather than the economics, that seems to worry Harris more.[11] (The argument is much the better for emphasizing the justice part. To point out that free riding reduces the likelihood of the research being done presupposes that it is important that it be done to begin with; but this is open to debate at the very least: I've already suggested that noone is wronged

by that research not being done. We might welcome further research, but unless we are already convinced that there is a generalized moral obligation that the research be done – which is far from self-evident – there being an incentive not to do it is not much of a cause for concern. By emphasizing the justice aspect, all that matters is that *someone* did *some* research from which we are benefiting without making a contribution of our own.)

Still: notwithstanding that the FRA is stronger than the other arguments examined so far won't show that it is powerful in its own right: those other arguments are decidedly weak. I'm willing to concede that there is at least some power in the FRA; but it is still not convincing. This is because the injustice that is claimed to be implicit in free riding is often not nearly as straightforward as the argument needs it to be.

Part of the problem is that the FRA helps itself to the justice claim rather too quickly. To see why, it's worth reconstructing a scenario in which there is free riding happening. Jack and Meg share a house. They agree that a house in which the kitchen is cleaned is better than one in which it is left uncleaned, and each is prepared to do the cleaning once the level of dirt on the hob has reached a certain level; however, each has a different dirt tolerance. Further, like anyone else, each would rather do as little housework as possible – and none at all if it can be helped. In this scenario, Meg has the lower dirt tolerance; therefore it will end up being her that does the cleaning, although the benefit of her efforts will be shared by both. Moreover, Jack and Meg are presented with an incentive to become more tolerant of dirt than the other: in this way, each knows that the kitchen will never become intolerable to them, but knows also that it will be the other person who will guarantee this (and might, with luck, keep the kitchen *pleasant* rather than minimally acceptable). In other words, each has an economic reason to become more slovenly than yesterday, since this means a high net gain in return for minimal work.

The problem in the Jack and Meg case is, at first glance, clear: Jack is free riding, this is unfair, and there is a straightforward moral reason not to act unfairly. But, on closer inspection, it's less clear that Jack really is as open to moral criticism as all that. He has not wronged Meg by not cleaning the kitchen. Had he not been there, she would have cleaned the kitchen anyway, and her appreciation of a clean kitchen ought in no way to be diminished by the fact that Jack appreciates it too: she would perhaps even be justified in taking pleasure in the fact that her activities had improved the welfare of others. Jack is a free rider because he capitalizes on what we might call 'superabundant' benefits; but the superabundance of these benefits is precisely why Jack is a free rider. We might, I suppose, want to say that free riding is disreputable, and that it would be admirable for Jack to lend more of a hand. But no harm or clear wrong is caused by being disreputable, so this will not show that he was under any obligation; a free rider's disrepute does not stem from his alleged injustice.

Chan and Harris have insisted that this claim is 'not strictly true', counterclaiming that Meg would not have had to clean so much if Jack had done his share:

> By such logic we could say that Jack would not wrong Meg by failing to pay his half of the rent, because under such circumstances she would most likely pay his share to avoid being evicted; but this is patently false![12]

This, though, is a misreading of my claim. The whole point of the example was that Meg *would* have done the cleaning anyway, subject to a proviso I shall outline in a moment, since she is less tolerant of dirt than Jack. Assuming that it is reasonable for people to tolerate *some* level of dirt – if they did not, they would never do anything but clean (maybe the morally enhanced people of Chapter 7 would be like this, but, as we've seen, their lives are not great) – the debate between Jack and Meg concerns what that tolerable level is, not whether he has or has not done his share once it is reached. It is less the case that Jack is taking Meg for a ride than that he is a straightforward beneficiary of her fastidiousness; she would have been no less fastidious without him around (and it seems very strange to say that Jack's obligations grow simply because of Meg's concerns). Unless it can be shown that Jack had committed some sort of ethical *faux pas* simply by being more easy-going in respect of crumbs, then there is nothing more to be said on the matter; Meg would (unless she is spiteful) do the cleaning either way. Conversely, if it can be shown that Jack had mis-stepped morally, then his free riding is not the problem: it's the mis-step that ought to concern us. Likewise, if there is an independent moral reason to be more supportive of research, then free riding in respect of its benefits seems not to be the primary concern.

In respect of Chan and Harris' rent counterexample, if the rent could only be paid because Jack and Meg contributed together and this was the understanding upon which they had signed the contract, then it would be wrong for him not to pay: in that case, Meg really would be exploited. But if he had just been allowed to sleep in the spare room, granted that it would have been admirable for him to offer to make a contribution towards board, the obligation is different. He is not committing an injustice by not doing so, and, the nature of a spare room being that it is *spare*, his not doing so does not undermine any of the benefits that Meg earned from paying the rent to begin with. The analogy doesn't fit; as such, it doesn't help the argument.

There is one proviso and one exception to this argument. I'll deal with the exception in the next section. The proviso is that the Jack and Meg example assumes that Jack does not contribute to the grubbiness of the kitchen. If he does leave spills that Meg invariably ends up cleaning, then his not helping to clean is less easily forgivable, since he is creating a burden for her. On the other hand, this obviously means that he is no longer *free* riding. And it's important to note that a person who does not assist or support research is not like this

anyway: while Jack, as it turns out, does make the kitchen a bit worse, and thereby increases the burden on Meg, the same cannot be said for people who do not support research. Although they do not make the world any better for others, neither do they create a burden. (Obviously, this point echoes the argument advanced against the PCA.) And if the objection to not participating in research is that it represents a failure positively to improve the world, the free-rider argument seems to have gone out of the window.

So: if it's true that people who don't support research are free-riders, it's not obvious that there's any injustice in that. Moreover, it's not obvious that people who benefit from others' research are free riding after all. This is because whatever ride one enjoys from scientific research does not, in fact, come for free. One pays in some way for that just about every medical benefit that one might enjoy. More than likely, if I live in the developed world, I have some form of medical insurance, publicly or privately provided, that covers interventions, some of them innovative, that restore or raise my level of flourishing to or towards what would be considered acceptable. If I'm interested in benefiting from interventions that make it possible for me to flourish at a more-than-minimally-reasonable level, the chances are that I would have to pay. Either way, the intervention is paid for; *de facto*, so is the research that makes that treatment possible. There is no sense in which I have acted as a free rider or broken elementary rules of justice. Therefore, if there *is* a duty to contribute to that which benefits me, it is one that I have already discharged; there is no further duty to research. It is possible that, in the case of therapeutic interventions, my insurance contributions are insufficient to pay for my treatment; but this is not too serious a point, since any contribution I might make individually to scientific progress would likely prove nugatory in terms of the effort that goes into important discoveries.

## Fairness and the future

Having said all this, there might be a narrow set of circumstances in which a fairness argument can be yoked to something that looks rather like a FRA, and it concerns those circumstances in which I am unwilling to contribute to the generation of some outcome that I desire and that I expect to accrue to me in the future anyway.

Imagine that concerts are put on only when musicians can be sure that enough people will buy tickets to cover the expenses. If I wait near the concert hall in the sure belief that other people will want and be willing to pay for a recital in sufficient numbers for it to happen, and if I do this with the intention of free riding by occupying an empty seat (I also believe that there will be one: the concert will be popular enough to be viable, but not a sell-out), then

I might well be acting unfairly. After all, I am expecting others to shoulder the burden of generating benefit on my behalf; what I am doing looks exploitative. Similarly, one might suppose that my refusal to support research when I fully intend to exploit the input of others is morally blameworthy, for reasons of fairness – noting, of course, that there is no appeal to be made here to benefits already received.

Nevertheless, if we do want to use an appeal to fairness and a consideration of future benefits as our moral motor, it is not clear that it would be all that powerful. For appeals to fairness will only stand a chance of getting anywhere when we do actually want and expect to benefit from the fruit of that research. Should we be indifferent, it would be hard to sustain an argument that we are acting unfairly by not supporting it: if anything, the unfairness charge could be more justly levelled at those who would have us carry a burden on their behalf for no clear benefit of our own. Should some research programme generate unexpected benefits – as may well be the case a good amount of the time – those benefits still will not generate the same obligation. In fact, it is hard to see how they generate any obligation at all.

For this reason, in effect, Chan and Harris' claim boils down to the instruction that we ought to pay for things that we want. True as this may be, it is also somewhat trivial. Moreover, while we might characterize the moral issue as one of fairness, it is not clear that we have to (at least, not when 'fairness' is meant to indicate 'recompense appropriate to benefits received'). We might feel that, by exploiting others, we wrong them – but while to wrong someone may be unfair, it doesn't follow that it is wrong *because* it is unfair. For example, we might think that the wrong has something to do with a neo-Lockean worry about the appropriation of another's labour and the fruits thereof; this being the case, the wrong has to do with undermining a person's rights. This makes fairness worries seem incidental, which looks like enough to show that there is no necessary link of the sort Chan and Harris seem to need between unfairness and wrongness.

What this means is that Chan and Harris have not established argumentatively that there is a generalizable duty to research based on appeals to fairness; and even if the case that has been made does imply some sort of duty in some circumstances, it is still not the same 'wide-angle' duty that they seem to think it is.

Naturally, this does not mean that we cannot still take on board some of what they propose. When they consider the extent of the duty to research that, ostensibly, we have, they indicate that we are required to carry our fair share of what is required across the board to support the research that benefits us. And this seems both correct and wholly compatible with the sort of neo-Lockean recasting of the obligation that we might imagine someone making: each of us is entitled to have certain expectations of others not to be exploited by them, and each of us may capitalize on the labour of others

to an extent limited by considerations of fairness and – for want of a better word – propriety. But, even here, it is also worth noting that, in this context, the fairness argument would have mutated slightly. No longer would we be saying that one ought to support research *because* it is fair; we would be saying that one ought to support research and that that duty (and others' duty to let us capitalize on that research) holds *to the extent* that is fair. In other words, fairness would merely describe the duty. It would not provide its origin.

## Reason and obligation

It is possible that what is at the core of my dispute with claims about there being a duty to research – and this speaks to the arguments in both this chapter and the last – is that their proponents are much too liable to treat a *reason* to act as an *obligation* to act. (This is a worry that is shared by Stuart Rennie.[13]) Rosamond Rhodes, for example, suggests that

> we would want medicine to be able to provide effective treatment when we or our loved ones should need it. Without human subject research, those treatments are less likely to be available. So, in light of our appreciation of human vulnerability to injury and disease and our appreciation of the value of clinical research, reasonable people should endorse policies that make research participation a social duty.[14]

She iterates the point elsewhere:

> We are all vulnerable to death, pain, disability, and the loss of pleasure and freedom that may be consequent to disease. These are conditions that we all would prefer to avoid, and when we or our loved ones are afflicted, we would all want a remedy to be available. Almost everyone has and almost all of their loved ones have medical needs at some point in their lives. Yet, we must acknowledge the need for improving on the standard of care for numerous conditions such as Alzheimer's disease, Parkinson's disease, sickle cell disease, end-stage renal disease, schizophrenia, and stroke. We need to learn about the causes and natural development of diseases, and the effectiveness of treatment for conditions such as burns, cancer, and severed spinal cords. The desired advance in treatment can only be achieved by studying our bodies.[15]

The essence of the claim is this: there are certain things that we would all want; therefore we are rationally required to endorse the means by which those things are realized. This kind of claim is echoed in the account of the obligation offered by Stjernschantz Forsberg et al.[16] Others give different reasons for wanting or endorsing further research; but the basic pattern of attempting to get those reasons to do the work of obligations is repeated elsewhere.

The example of the sort of person who prefers artistic pursuit to the pursuit of cures even for illnesses that afflict him in the last chapter should undermine

confidence in claims about what we would all prefer, since there's nothing irrational about a given set of preferences; but even putting that quibble aside, it should be clear that there's a problem with the leap being made: though there may be certain pro-attitudes that are rationally required in some contexts – it might not be rational to endorse a set of ends (such as the fruits of research) without endorsing the means (such as research) necessary to realize them – what is rationally required and what is morally required are not the same thing. Hence establishing that we have a reason to endorse something like research is not going to be enough to establish that there is any other kind of requirement. For sure, an obligation is a reason to act; but rational requirements are not moral requirements – and certainly not simply by virtue of being rationally required. So even if we grant that beneficence, generation of public good (which Schaefer et al. consider to be different from beneficence[17]), rescue, filial piety, harm prevention, fairness, and all the rest of it are *reasons* to research – or even moral reasons – we are still not committed to thinking that we're faced with an obligation.

We can apply the excavation of the fallacious conflation of having reasons and having obligations back to Chan and Harris' argument about free riding. In their 2009 paper, they claim that 'if free riding undermines a social practice [such as scientific research] that is by hypothesis morally and practically important, then it is not rational'[18]. This statement is perplexing, and raises several questions. Is it true? Why should it be irrational to undermine even something that is morally important? And how does a claim about irrationality have any moral weight? It might be immoral to undermine valuable institutions, but rationality has nothing obvious to do with it. If rationality *is* important, what makes it *morally* important?

It seems not to be true that it is always irrational to act in a way that undermines something 'morally and practically important'. To see why, assume that a functional banking system is at least practically important, and imagine that someone believes that this system is about to collapse. Motivated by this belief, he rushes to withdraw all his savings. This seems to be a rational thing to do. However, because only a minority of bank assets are liquid, such a withdrawal will actually make the atrophy of the banking system a little more likely. In other words, this is an example of a rational action undermining an institution that is notionally practically important; and so we have falsified the idea that undermining such institutions is necessarily irrational. It won't do to insist that one person's withdrawing his savings is unlikely to make a difference in the grand scheme of things, because neither will one person's participation or non-participation in research. Besides: if this riposte won't do, we can just imagine that many people share a fear about the banking system, and each individually withdraws his savings. This would amount to a run on the banks – and runs on banks do damage them. Yet the behaviour is still not irrational – indeed, if you suspect that a run on the bank is beginning,

withdrawing your savings as soon as possible might be the rational thing to do, notwithstanding that people withdrawing their savings too quickly is at the root of the immediate problem.

We might wonder whether there's a wedge to be driven between things that are morally important and things that are practically important: maybe the banking analogy speaks to the latter, but not the former, so that it could still be irrational to undermine things that are morally important. Let's allow that that's a possibility: it doesn't matter, though, because – in an echo of the PCA – a failure to participate in research, or even endorse it, doesn't undermine any established practice anyway. All it does is fail to promote it.

What free riding might do, admittedly, is undermine the possibility that a practice that is not established will get going. But, as I indicated above, this is an economic problem; it is not self-evident that we have to treat it as a moral problem – or that we would be entitled to treat it as such in the absence of either a supplementary claim about the importance of establishing a practice such as research (which would be question-begging), or a claim that what is irrational is therefore immoral. But that supposition is shaky simply because, as I've said, the requirements of reason and the requirements of morality are different kinds of thing. To hold that establishing irrationality is *de facto* immorality would commit us to the belief that foolishness is a sin – and that's implausible.

## A puzzle about duties

One more thing is worth noting. Let's pretend for the moment that the FRA and the AR are successful. What is puzzling is that the two arguments are yoked together, since they seem to be antagonistic. For a person who is rescued is, in a sense, the *excelsior* of free riders: he is a person who receives what I shall allow to be the greatest possible benefit from a second person who, it would appear, receives nothing – or nothing comparable – in return. While the person rescued might owe a debt of thanks to the rescuer, and we might think him pusillanimous if he offered no token of appreciation, we would think it odd, if not monstrous, for the rescuer to insist on some sort of recompense. We might even think that the rescuer has a duty of some sort at least to act as though even gratitude was wholly unnecessary. At the same time, we might think that a sincere offer of thanks offsets whatever moral debt one might have accrued in being rescued; but still, if the rescued offered no thanks to the rescuer, to point out that he was acting as a free rider would be a curiously etiolated kind of criticism. Rescuers have a duty to treat the rescued as *unproblematic* free riders.

Now, the puzzle is this: Harris wants to be able to say that past research puts me in a moral debt that I have a duty to repay, and that not to repay this moral debt is unjust. He is happy to claim that merely coming into existence is in a person's interest[19] – in other words, that he is better off existing than not. So

far, so good. But I do not see why a person could not use this claim to suggest that, given that he owes his existence to a combination of factors including past research, this research has in effect rescued him – whether from non-existence or from a worse existence than he currently has does not matter, and nor does it matter that he needs someone like Harris to point out that he was rescued. He might then feel entitled to deny that he owes anything beyond thanks. What is going on here?

We might respond by looking at the motivation of the person who did the research from which our subject is currently benefiting. Doing so will reveal that there is a serious, possibly fatal, tension in Harris' argument. If the motivation of the researcher was to rescue persons from non-existence or seriously less-good lives, there can be no cause for complaint if they recognize no further duty on them to research, for just the same reason that a person rescued from a pond might have a duty to be grateful (although gratitude expressed for the sake of duty is, of course, no such thing) but no concomitant duty beyond that. But if the person who benefited them was motivated by the sting of the FRA or the argument from filial piety, those who benefit *still* have no duty to participate in further research. Indeed, it seems possible to me that the researcher who benefited them might owe *them* and everyone else who has benefited some kind of gratitude for having afforded him the chance to pay off the social debt that *he* feels so acutely.

Put bluntly, if the research from which I am currently benefiting was motivated by claims about justice, I can easily say that I am not free riding so much as enjoying what was due to me as a member of the society that supported the research in the first place. If it was motivated by rescue, the idea that commensurate obligations befall me in return goes out of the window. I think that I am entitled to go from here to quite a bold claim on this basis: for as long as the person the fruits of whose research are available to me acted *for a reason*, I am released from any duty to research.

It is on this basis that the mere fact that I have benefited from a charity's funding of research imposes no duty to research on me: if those charity workers were motivated by a sense that they had to rescue me from cardiac disease, I might owe them hearty gratitude, but I am not obliged to further the cause of cardiology any more than a rescued child is obliged to promote water safety. If the charity's workers were motivated by a desire to expiate some kind of moral debt accrued by their having benefited from things like research, there is also no obligation on me – to deny this would be like saying that a child conceived out of his parents' sense of duty to the family line is, *for that reason alone*, obliged to beget children of his own. It would be absurd.

Strictly speaking, none of what I have said in this chapter or the last will prove definitively that there is no duty to research. However, none of what I take to be the strongest arguments in favour of there being such a duty is convincing. Where obligations are generated, they are not reliably the *right* obligations.

And even if we think that the AR and the FRA are successful in their own terms, it is worth noting that we ought to be wary of making the arguments too closely together, since the former does seem rather to undermine the foundations of the latter. If it is possible to provide answers to questions about the motivation of past researchers that ascribe to them reasons for doing what they did, the idea that *I* might have Harris-like reasons to research evaporates.

Research may yield things that would contribute significantly to the flourishing of others; and, inasmuch as that it's an admirable thing to do, and doing admirable things is plausibly a sign of a life well-lived, it might indicate a kind of flourishing on the part of the agents doing the research. But, still, it falls a long way short of being obligatory.

## Notes

**1** Harris and Holm, 2003, pp. 123–4.
**2** Harris, 2007, p. 189.
**3** Chan and Harris, 2009, p. 165.
**4** Harris, 2005, p. 242.
**5** Ibid., p. 245.
**6** Singer, 1996, pp. 229ff; cf. Singer, 1972, p. 231.
**7** cf. Rennie 2011, p. 41: 'if an action is morally obligatory for a person, then that person is the appropriate object of negative moral attitudes (such as blame) if he or she fails to perform that action, unless he or she has a legitimate excuse', though it's not clear whether we're supposed to think that something is obligatory because failure to perform it attracts appropriate disapproval, or whether failure to perform it attracts that disapproval because it's obligatory.
**8** Chan and Harris, 2009, *passim*, but esp. pp. 162–3.
**9** cf. Shapsay and Pimple, 2007, *passim*.
**10** Harris, 2005, p. 243.
**11** Ibid., p. 243.
**12** Chan and Harris, 2009, p. 164.
**13** Rennie, 2011, p. 42.
**14** Rhodes, 2005, p. 15.
**15** Rhodes, 2008, p. 37.
**16** Stjernschantz Forsberg et al. 2013, *passim*.
**17** vide Schaefer et al. 2009, passim; Rennie, 2011, p. 42.
**18** Chan and Harris 2009, p. 164.
**19** Harris, 2004, p. 70.

# 9-and-a-bit

# Bioscience and the Good Life

Debates about the use of the biosciences and enhancement technologies can have a tendency to be dominated by optimistic enthusiasts on one hand, telling us how we should embrace new technologies (perhaps remaking human nature in the process), and more pessimistic conservatives on the other, warning us about the moral pitfalls into which we could too-easily stumble. I hope that this book has, on the whole, managed to avoid those poles. I've tried to make the case that there is nothing wrong with the use of most enhancement technologies, but that we should not get carried away by their potential, and that there is certainly no obligation to embrace them. The moral reasons that we have to endorse therapy – the kind of intervention that aims to raise a life, or an aspect of a life, to a level of flourishing at least minimally acceptable – do not beget reasons to endorse enhancement. And where lives can be made better, perhaps by technological means, it doesn't follow that it has to be by biotechnological means. The kinds of silicone-based cognitive enhancer that we carry around in our pockets obviate the need for cognitive enhancers that we carry around inside our skulls, for example. In the last two chapters, I've tried to push this claim a little further, to deny that there is any particular obligation to support scientific research. A significant part of that is the hypothesis that, while there is no obligation to improve already-good-enough lives, even where there is an obligation or desire to improve lives, we can do so in ways that do not require biotechnological innovation.

The first three chapters were essentially directed at clearing the path for the next four. I took as my starting point the idea that ethics concerns itself with the examination of human behaviour. Sometimes this examination is concerned with whether a particular practice is right or wrong, and with the standard by which we're supposed to judge that. But Aristotle points to another role for the ethicist, which is to ask whether, granted certain claims about the end to which an activity is directed, that activity is really going to help us realize it. For Aristotle, all roads lead to the same end: human action is directed (admittedly, sometimes indirectly) at securing the good life – something that is captured by the idea of flourishing or *eudaimonia*. The potential problem with this is that what constitutes the good life is only vaguely defined (there's a nagging worry that there's nothing more to it than the circular 'whatever it is at which all action is directed'), and there's lots of scope for conflict about the means to realize it. Yet if one is suspicious of this kind of Aristotelian picture, in

which we have to accept that securing the good life is the *sole* end of human activity, that isn't necessarily too big a problem: one can still take seriously the weakened idea that it is one of the things at which the biosciences are designed to secure. Either way, if one of the things that ethics can do is to help work out whether or not some activity is conducive to the good life, then it follows that ethics can have something to say about the biosciences – and something that is potentially richer than a set of statements about what is and is not permissible.

An advantage of this approach is that it allows for a narrower and hopefully more focused ethical appraisal of the biosciences to be taken than might otherwise be possible. Given the range of bioscientific endeavour – from deepening our understanding of cancer genetics to synthetic biology – the range of things that might warrant questioning is vast, as is the range of ways in which they might be questioned. (Questions about permissibility would invite higher-order questions about how that permissibility is to be measured, for example: are we consequentialist or non-consequentialist, for starters.) Taking it as read that one of the things that the biosciences is supposed to do is to improve the lives of humans, and asking whether that aim is likely to be achieved, means not having to worry about any of that.

Still, questions about permissibility are hard to avoid entirely. If the applications of the biosciences at which I was planning to look – enhancements in particular – were impermissible from the start, there would be nothing more to say. But, as we've seen from Chapter 2, arguments that enhancements are impermissible are very hard to make work. At the simplest level, to be against enhancement means to be against therapy, since therapy is a kind of enhancement. But even if – as I think we should – we think that there's a tenable distinction to be drawn between therapy and enhancement, and interpret the anti-enhancement arguments as being against interventions that are directed not at putting right something that's gone wrong, but at 'supernormal' improvements, the arguments don't amount to all that much.

I am satisfied that there is no obligation not to enhance. But neither is there an obligation to do so, as I argued in Chapter 3. Partly this is because 'enhance' is just too broad a category. But it's also because there would only be a duty to enhance if it could be shown that it would be blameable not to. It's unlikely that this can be shown, at least when 'enhancement' is treated as indicating something non-therapeutic. Enhancement, properly so-called, would refer to interventions beyond what is required to make a life minimally acceptable. But establishing that there is a duty of beneficence that manifests in this way is harder than it looks – and the reasons for this played a significant role in the argument of Chapter 8, as well. Bluntly, a duty to enhance amounts to a duty to make the world better than that which reasonable people would agree to be acceptably good. I do not see any reason to suppose that there is such a duty.

The moral status of enhancements is that they're permissible, but nothing more. With that in mind, I turned to the question of whether, and to what extent, they might contribute to flourishing or the good life – whether there's much of a reason short of a duty to pursue them. Though it's impossible to talk in much detail about every possible enhancement, by looking at enhancement technology in the broadest terms, it's been possible to say that it's unlikely that there is. Few would contribute substantially to the good life, if at all.

The prospect of indefinite life extension – which I take to be the paradigmatic enhancement – turns out to be less desirable than we might instinctively think. The reason for that is derived from the idea (which I think we should accept) that being alive is not its own reward, but is rather the foundation upon which the edifice of a life is built – and it's the edifice that matters. This, of course, takes us back to the 'projects' claim from Chapter 1. It's *doing* these things that counts, though, more than having the opportunity to have done them. It's being engrossed in a project or activity that is desirable. It is, I think, better to go out while doing something fulfilling, than to have to find something else to make life fulfilling. What should concern us is that we are living a good life when we die; and dying later will not make that life any better. Indeed, worrying about the proximity of death might itself erode the quality of the life being lived.

Perhaps counterintuitively, I think that the biosciences are more likely to make a meaningful contribution to a good life when they are used in the pursuit of things that we might think less essential. In Chapter 5, I looked at the possibility of making alterations to the body for the sake of sporting ambition, and for the sake of artistic endeavour; but the principle extends more generally. Once we've accepted, first, that the extension of life for life's own sake is not all that desirable a use of the biosciences, but, second, that the pursuit of projects of some sort is a significant consideration in what makes a life good, and that opens the way to seeing what one of the roles of the biosciences could be for making life better.

None of these projects has (or need have) any particular moral weight of their own: all that matters is that a person might adopt such-and-such a project, and treat the opportunities afforded by developments in science and technology as tools in the realization of that project. However, I am more confident that this assertion would be truer in respect of some enhancements than others. Notably, I think that it is much less likely to be true in respect of cognitive enhancers – again, somewhat in contrast to some of the more convinced of the optimists.

Not all enhancements are pursued for the sake of the enhancee's flourishing. Something like moral enhancement, and the idea that we might be able to take advantage of knowledge about the neurological, neurochemical and genetic correlates of propensities towards desirable and undesirable behaviour, fits into this category. Yet there's a number of problems with an overhasty

adoption of moral enhancement. The big problem with moral enhancement is that it runs the risk of being much too certain about what traits are desirable and ought to be promoted; but if there's a reasonable chance of fallibility here, we run the risk of making people less morally lithe than they would need to be to reboot and refresh their moral proclivities. Maintaining the distinction between acting *for* and *from* a cause keeps the possibility of such a reboot higher.

And even if we bite the bullet and say that what we're interested in is not so much *moral* enhancement as heightened pro-sociality, this might lead to independent moral objections of three sorts. One would be that we stand in danger of preferring to limit some people's moral dignity for the sake of our – fallible – ideas about what it would be desirable for them to be like; and if we think that moral dignity is a component of the good life and the promotion of the good life interests us, there is obviously a tension here. Another is that even if certain characteristics are *ceteris paribus* reliably vicious, it is not a given that we would want to eliminate all vice from the world anyway: a uniformly virtuous world might be worse in important (though perhaps non-moral) respects than we would expect, or even find particularly desirable. And, finally, once again, we'd have departed from the idea of using the biosciences to improve the lives of the people on whom they're used, and turned to the idea of using them for the sake of improving the lives of others. This may or may not be acceptable – but it's something that requires justification of its own.

The provisional conclusion reached by the end of Chapter 7, then, was this: that of the ways that we might imagine that the biosciences and the applications that arise from them could improve flourishing and make lives better, a good number are not, after all, so straightforwardly desirable. Assuming that the people who would benefit from enhancement technologies already had lives that were at least minimally acceptably good, a good number of those technologies may do little to improve their lives further – or, at the very least, will do little that could not be matched by other, non-bioscientific means.

My overarching conclusion, then, runs along these lines: though it's true that the biosciences and the technologies they've facilitated have made our lives better than they might have been in some other world where those advances had not been made, it doesn't follow from that that bioscientific advance is a necessary part of the good life, or of making lives better. We ought not to shy away from their use; but neither is there any particular need to expend a huge amount of energy realizing bioscientific promises. And if this is right, then there doesn't seem to be much to underpin the idea that there's a duty to research more of them. The only rationale for further research would be therapeutic. Improving the therapeutic repertoire may be admirable: but it is not obligatory, though; and what obligations there are need not be met by means of bioscientific innovation.

I began this book with an epigram taken from Olaf Stapledon's *Last and First Men*, and though the book is cranky (and, frankly, quite boring), it's with another quotation – drawn from a couple of paragraphs before the first one – that I think I'll draw it to a close:

> It was such a society, almost unthinkable to the First Men, that now set about remaking human nature. Unfortunately there were conflicting views about the goal.[1]

It is not inconceivable that a philosopher may have helped these future humans decide on the goal to be pursued. But neither is it certain that they'd have agreed that there was any real need to alter or enhance humanity in any way at all.

## Note

**1** Stapledon, 2004, p. 186.

# Bibliography


Adler, R., (2010), 'Ray Kurzweil: Building bridges to immortality'. *New Scientist* 2792: 38–9.

Agar, N., (2010), *Humanity's End* (Cambridge, MA: MIT Press).

— (2013), 'Why is it Possible to Enhance Moral Status and why is doing so Wrong?'. *Journal of Medical Ethics* 39(2): 67–74.

Anon, (2009a), 'Phelps Coach Wades into Suit Row', BBC Sport, via http://www.news.bbc.co.uk/sport1/hi/other_sports/swimming/8172910.stm.

— (2009b), 'Boys' mother "fed them cannabis"', BBC News online, via http://www.news.bbc.co.uk/1/hi/england/south_yorkshire/8235661.stm.

— (2010a), 'Girl to have Leg Amputated to Become Paralympian'. *Manchester Evening News*, via http://www.menmedia.co.uk/manchestereveningnews/news/s/1301780_girl_to_have_leg_amputated_to_become_paralympian.

— (2010b), 'Edlington attack boys had "toxic" upbringing', BBC News online, via http://www.news.bbc.co.uk/1/hi/england/south_yorkshire/8473112.stm.

Anscombe, G. E. M., (1958), 'Modern Moral Philosophy'. *Philosophy* 33(124): 1–19.

Arendt, H., (1994), *Eichmann in Jerusalem: A Report on the Banality of Evil* (New York: Penguin).

Aristotle, (1976), *Ethics* (London: Penguin).

— (1992), *Politics* (London: Penguin).

Austin, M. W., (2009), 'Magnanimity, Athletic Excellence, and Performance-Enhancing Drugs'. *Journal of Applied Philosophy* 26(1): 46–53.

Bakewell, S., (2011), *How to Live: A Life of Montaigne in one Question and Twenty Attempts at an Answer* (London: Vintage).

Ball, P., (2007), *The Devil's Doctor: Paracelsus and the World of Renaissance Magic and Science* (London: Arrow).

Balazs, A. B., Chen, J., Hong, C. M., Rao, D. S., Yang, L. and Baltimore, D., (2012), 'Antibody-Based Protection against HIV Infection by Vectored Immunoprophylaxis'. *Nature* 481(7379): 81–4.

Beauchamp, T. and Childress, J., (2009), *Principles of Biomedical Ethics, 6th edn.* (New York: Oxford University Press).

Beckett, S., (2006), *Waiting for Godot* (London: Faber).

Bennett, R., (2009), 'The Fallacy of the Principle of Procreative Beneficence'. *Bioethics* 23(5): 265–73.

Bernstein, D., et al., (2005a), 'False Memories about Food can lead to Food Avoidance'. *Social Cognition* 23(1): 11–34.

— (2005b), 'False Beliefs about Fattening Foods can have Healthy Consequences'. *Proceedings of the National Academy of Sciences* 102(39): 13727.

Borges, J. L., (1998), *Fictions* (London: Calder).

Borglin, G., Edberg, A. K. and Hallberg, I. R., (2005), 'The Experience of Quality of Life among Older People'. *Journal of Aging Studies* 19(2): 201–20.

Bostrom, N., (2008), 'Smart Policy: Cognitive Enhancement in the Public Interest', in Zonneveld, L., Dijstelbloem, H., and Ringoir, D., (eds.) *Reshaping the Human Condition: Exploring Human Enhancement* (The Hague: Rathenau Institute, in collaboration with the British Embassy, Science and Innovation Network and the Parliamentary Office of Science & Technology).

Bostrom, N. and Ord, T., (2006), 'The Reversal Test: Eliminating the Status Quo Bias in Applied Ethics'. *Ethics* 116(4): 656–79.

Bostrom, N. and Sandberg, A., (2009), 'Cognitive Enhancement: Methods, Ethics, Regulatory Challenges'. *Science and Engineering Ethics* 15(3): 311–41.

Brassington, I., (2007), 'On Heidegger, Medicine, and the Modernity of Modern Medical Technology'. *Medicine, Health Care, and Philosophy* 10(2): 185–95.

— (2013a), 'What's the Point of Philosophical Bioethics?'. *Health Care Analysis* 21(1): 20–30.

— (2013b), 'The Case for a Duty to Research: Not Yet Proven'. *Journal of Medical Ethics*, doi: 10.1136/medethics-2013-101370.

British Medical Association, (2004), 'The Law and Ethics of Male Circumcision'. *Journal of Medical Ethics* 30(3): 259–63.

Buchanan, A., (2009), 'Moral Status and Human Enhancement'. *Philosophy and Public Affairs* 37(4): 346–81.

— (2011), *Beyond Humanity?* (Oxford: Oxford University Press).

Buckholtz, J. W., Asplund, C. L., Dux, P. E., Zald, D. H., Gore, J. C., Jones, O. D. and Marois, R., (2008), 'The Neural Correlates of Third-Party Punishment'. *Neuron* 60(5): 230–40.

Burgess, A., (1972), *A Clockwork Orange* (London: Penguin).

Carel, H., (2009), '"I am Well, Apart from the Fact that I have Cancer": Explaining Well-being within Illness', in Bortolotti, L. (ed.), *Philosophy and Happiness* (London: Palgrave Macmillan).

— (2011), *Illness* (Durham: Acumen).

Carroll, L., (1998), *Alice's Adventures in Wonderland/Through the Looking-Glass* (Oxford: Oxford World's Classics).

Chadwick, R., (2008), 'Therapy, Enhancement and Improvement', in Gordijn, B. and Chadwick, R. (eds), Medical Enhancement and Posthumanity (Dordrecht: Springer).

Chan, S. and Harris, J., (2009), 'Free Riders and Pious Sons – Why Scientific Research Remains Obligatory'. *Bioethics* 23(3): 161–71.

— (2011), 'Moral Enhancement and Pro-Social Behaviour'. *Journal of Medical Ethics* 37(3): 130–1.

Cochrane, A., (2010), 'Undignified Bioethics'. *Bioethics* 24(5): 234–41.

Conan Doyle, A., (2009), 'A Scandal in Bohemia', repr.in *The Penguin Complete Sherlock Holmes* (London: Penguin, 2009).

Crockett, M., et al., (2010a), 'Serotonin Selectively Influences Moral Judgment and Behavior through Effects on Harm Aversion'. *PNAS* 107(40): 17433–8.

— (2010b), 'Reply to Harris and Chan: Moral Judgment is more than Rational Deliberation'. *PNAS* 107(50): E184.

DeGrazia, D., (2013), 'Moral Enhancement, Freedom, and what we (should) Value in Moral Behaviour'. *Journal of Medical Ethics,* doi:10.1136/medethics-2012-101157.

de Grey, A., (2004), 'Leon Kass: Quite Substantially Right'. *Rejuvenation Research* 7(2): 89–91.
— (2005), 'Life Extension, Human Rights, and the Rational Refinement of Repugnance'. *Journal of Medical Ethics* 31(11):659–63.
den Hartogh, G., (2001), 'The Slippery Slope Argument', in Kuhse, H. and Singer, P. (eds), *A Companion to Bioethics* (Malden, MA: Blackwell, 2001), pp. 280–90.
Douglas, T., (2008), 'Moral Enhancement'. *Journal of Applied Philosophy* 25(3): 228–45.
— (2011), 'Moral Enhancement via Direct Emotion Modulation: A Reply to John Harris', doi: 10.1111/j.1467-8519.2011.01919.x.
— (2013), 'The Harms of Status Enhancement could be Compensated or Outweighed: A Response to Agar'. *Journal of Medical Ethics* 39(2): 75–6.
Everett, D., (2009), *Don't Sleep, There are Snakes* (London: Profile).
Farah, M, et al., (2004), 'Neurocognitive Enhancement: What Can we do and what Should we do?'. *Nature Reviews Neuroscience* 5(May): 421–5.
Ferguson, R. and Chang, D., (2010), 'Postmastectomy Breast Reconstruction', in Siemionow, M. and Eisenmann-Klein, M. (eds), *Plastic and Reconstructive Surgery* (London: Springer).
Ferris, J., (2008), *Then We Came to the End* (London: Penguin).
Fine, C., (2011), *Delusions of Gender: The Real Science behind Sex Differences* (London: Icon).
Fineman, H., (1994), 'An Older, Grimmer Jesse'. *Newsweek* 123(2): 24.
Fukuyama, F., (2002), *Our Posthuman Future* (New York: Farrar, Straus and Giroux).
Gabriel, Z. and Bowling, A., (2004), 'Quality of Life from the Perspectives of Older People'. *Ageing & Society* 24(5): 675–91.
Gould, E., (2009), 'Childhood Lead Poisoning: Conservative Estimates of the Social and Economic Benefits of Lead Hazard Control'. *Environmental Health Perspectives* 117(7): 1162–7.
Greene, G. and Reed, C., (1968), *The Third Man* (London: Lorrimer).
Habermas, J., (2008), *The Future of Human Nature* (Cambridge: Polity).
Harris, J., (2001), 'One Principle and Three Fallacies of Disability Studies'. *Journal of Medical Ethics* 27(6): 383–7.
— (2004), *On Cloning* (London: Routledge).
— (2005), 'Scientific Research is a Moral Duty'. *Journal of Medical Ethics* 31(4): 242–8.
— (2007), *Enhancing Evolution* (Princeton: Princeton University of Press).
— (2011), 'Moral Enhancement and Freedom'. *Bioethics* 25(2): 102–11.
— (2012a), 'What it's Like to be Good'. *Cambridge Quarterly of Healthcare Ethics* 21(3): 293–305.
— (2012b), 'Moral Progress and Moral Enhancement'. *Bioethics,* doi:10.1111/j.1467-8519.2012.01965.x.
— (2012c), '"Ethics is for Bad Guys!": Putting the "Moral" into "Moral Enhancement"'. *Bioethics,* doi:10.1111/j.1467-8519.2011.01946.x.
— 'Taking Liberties with Free Fall'. *Journal of Medical Ethics,* doi: 10.1136/medethics-2012-101092.

Harris, J. and Chan, S., (2010), 'Moral Behavior is not what it Seems'. *PNAS* 107(50): E183.
Harris, J. and Holm, S., (2003), 'Should We Presume Moral Turpitude In Our Children? Small Children And Consent To Medical Research'. *Theoretical Medicine* 24(2): 121–9.
Hauskeller, M., (2013), 'The moral status of post-persons'. *Journal of Medical Ethics* 39(2): 76–7.
Havelock, E., (1982), *Preface to Plato* (Cambridge, Mass and London: Harvard University Press).
Heidegger, M., (1999), *Basic Writings* (London: Routledge).
Herissone-Kelly, P., (2006), 'Procreative Beneficence and the Prospective Parent'. *Journal of Medical Ethics* 32(3): 166–9.
— (2012), 'Wrongs, Preferences and the Selection of Children: A Critique of Rebecca Bennett's Argument against the Principle of Procreative Beneficence'. *Bioethics* 26(8): 447–54.
Hodge, J., (1995), *Heidegger and Ethics* (London: Routledge).
Höffe, O., (2010), *Can Virtue make us Happy?* (Evanston, IL: Northwestern University Press).
Homer, (1991), *Odyssey* (Harmondsworth: Penguin).
Horgan, J., (2005), 'Cultured bone offers novel wedding rings'. *New Scientist* 2488: 15.
Hume, D., (1985), *A Treatise of Human Nature* (Harmondsworth: Penguin).
Huxley, A., (1977), *Brave New World* (London: Grafton).
Jensen, D., (2011), 'A Kantian Argument against Comparatively Advantageous Genetic Modification'. *Journal of Medical Ethics* 37(8): 479–82.
Johnson, B. S., (2001), *Christie Malry's Own Double-Entry* (London: Picador).
Jonas, H., (1969), 'Philosophical Reflections on Experimenting with Human Subjects'. *Daedalus* 98(2): 219–47.
Kant, I., (1993), *Grounding for the Metaphysics of Morals* (Indianapolis: Hackett).
— (1998), *The Metaphysics of Morals* (Cambridge: Cambridge University Press).
— (2004), *Religion within the Boundaries of Mere Reason* (Cambridge: Cambridge University Press).
Kass, L., (1997), 'The Wisdom of Repugnance'. *Valparaiso University Law Review* 32(2):679–706.
— 'Ageless Bodies, Happy Souls: Biotechnology and the Pursuit of Perfection'. *The New Atlantis* Number 1 (Spring 2003): 9–28.
Kurzweil, R. and Grossman, T., (2011), *Transcend: Nine Step to Living Well Forever* (New York: Rodale).
Loftus, E. and Berenstein, D., (2005), 'Rich False Memories: The Royal Road to Success', in Healy, A. (ed.), *Experimental Cognitive Psychology and its Applications* (Washington DC: American Psychological Association Press).
Lyons, B., (2012), 'Solidarity, Children and Research'. *Bioethics* 26(7): 369–75.
Mackie, J. L., (1977), *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin).

Maguire, E. A., Gadian, D. G., Johnsrude, I. S., Good, C. D., Ashburner, J., Frackowiak, R. S. J. and Frith, C. D., (2000), 'Navigation-Related Structural Change in the Hippocampi of Taxi Drivers'. *PNAS* 97(8): 4398–403.
— (2006), 'London Taxi Drivers and Bus Drivers: A Structural MRI and Neuropsychological Analysis'. *Hippocampus* 16(12): 1091–101.
Mandeville, B., (1997), *The Fable of the Bees and Other Writings* (Indianapolis: Hackett).
Marín-Gutzke, M. and Sánchez-Olaso, A., (2010), 'Reconstructive Surgery in Young Women with Breast Cancer'. *Breast Cancer Research and Treatment* 123(Supp 1):67–74.
Marx, K., (1970), *Economic and Social Manuscripts of 1844* (London: Lawrence and Wishart).
McDowell, J., (1998), *Mind, Value and Reality* (Cambridge, MA: Harvard University Press).
McKeegan, A., (2010), 'Amputee Girl, 11, with Paralympic Dream Takes Giant Strides'. *Manchester Evening News*, via http://www.menmedia.co.uk/manchestereveningnews/news/s/1314170_amputee_girl_11_with_paralympic_dream_takes_giant_strides.
Meyer-Schönberger, V., (2009), *Delete: The Virtue of Forgetting in a Digital Age* (Woodstock: Princeton University Press).
Mielke, H. and Zahran, S., (2012), 'The Urban Rise and Fall of Air Lead (Pb) and the Latent Surge and Retreat of Societal Violence'. *Environment International* 43(August): 48–55.
Milgram, S., (1974), *Obedience to Authority* (London: Tavistock).
Mill, J. S., (1874), *Three Essays on Religion* (London: Longmans, Green, Reader, and Dyer).
— (1998), 'Utilitarianism', repr., in John Gray (ed.), *On Liberty and Other Essays* (Oxford: Oxford World's Classics).
Millgram, E., (2005), *Ethics Done Right* (Cambridge: Cambridge University Press).
Mithoefer, M., et al., (2013), 'Durability of Improvement in Post-Traumatic Stress Disorder Symptoms and Absence of Harmful Effects or Drug Dependency after 3,4-Methylenedioxymethamphetamine-Aassisted Psychotherapy: a Prospective Long-term Follow-up Study'. *Journal of Psychopharmacology* 27(1): 28–39.
Montaigne, M., (1991), *The Complete Essays* (London: Penguin).
Morrison, T., (1997), *Beloved* (London: Vintage).
Nagel, T., (1970), 'Death'. *Noûs* 4(1): 73–80.
Nevin, R., (2007), 'Understanding International Crime Trends: The Legacy of Preschool Lead Exposure'. *Environmental Research* 104(3): 315–36.
Nietzsche, F., (1974), *The Gay Science* (New York: Vintage).
— (1992), *Ecce Homo* (Harmondsworth: Penguin).
Nussbaum, M., (1996), *The Therapy of Desire* (Princeton: Princeton University Press).
— (2006), *Frontiers of Justice* (Cambridge, MA: Belknap).
Omohundro, S., (2009), 'AI and the Future of Human Morality', via http://www.acceleratingfuture.com/people-blog/2009/ai-and-the-future-of-human-morality/.

Oultram, S., 'All Hail the new Flesh: Some Thoughts on Scarification, Children and Adults'. *Journal of Medical Ethics* 35(10): 607–10.
Ovid, (1982), *The Erotic Poems* (London: Penguin).
Parfit, D., (1987), *Reasons and Persons* (Oxford: Clarendon).
Parker, E., et al., (2006), 'A Case of Unusual Autobiographical Remembering'. *Neurocase* 12(1): 35–49.
Persson, I. and Savulescu, J., (2008), 'The Perils of Cognitive Enhancement and the Urgent Imperative to Enhance the Moral Character of Humanity'. *Journal of Applied Philosophy* 25(3): 162–77.
— (2012), *Unfit for the Future: The Need for Moral Enhancement* (Oxford: Oxford University Press).
Plato, (1952), 'Phaedrus', in Hackforth, R (ed.), *Plato's Phaedrus* (Cambridge: Cambridge University Press).
— (1981), 'Crito', repr., in *Five Dialogues* (Indianapolis: Hackett).
— (1999), *Theaetetus* (Oxford: Clarendon).
Ramachandra, M., et al., (2001), 'Re-Engineering Adenovirus Regulatory Pathways to Enhance Oncolytic Specificity and Efficacy'. *Nature Biotechnology* 19(11): 1035–41.
Rennie, S., (2011), 'Viewing Research Participation as a Moral Obligation: In Whose Interests?'. *Hastings Center Report* 41(2): 40–7.
Rhodes, R., (2005), 'Rethinking Research Ethics'. *American Journal of Bioethics* 5(1): 7–28.
— (2008), 'In Defense of the Duty to Participate in Biomedical Research'. *American Journal of Bioethics* 8(10): 37–8.
Rosoff, P., (2012), 'The Myth of Genetic Enhancement'. *Theoretical Medicine and Bioethics* 33(3): 163–78.
Ross, W. D., (2002), *The Right and the Good* (Oxford: Clarendon).
Sandel, M., (2007), *The Case against Perfection* (Cambridge, MA: Belknap).
Savulescu, J., (2001), 'Procreative Beneficence: Why We should Select the Best Children'. *Bioethics* 15(1–6): 413–26.
— (2010), 'Why Bioenhancement of Mathematical Ability Is Ethically Important', *Practical Ethics* blog, via http://www.blog.practicalethics.ox.ac.uk/2010/11/why-bioenhancement-of-mathematical-ability-is-ethically-important/.
Savulescu, J. and Kahane, G., (2009), 'The Moral Obligation to Create Children with the Best Chance of the Best Life'. *Bioethics* 23(5): 274–90.
Savulescu, J., et al., (2006), 'Behavioural Genetics: Why Eugenic Selection is Preferable to Enhancement'. *Journal of Applied Philosophy* 23(2): 157–72.
Schaefer, O., et al., (2009), 'The Obligation to Participate in Biomedical Research'. *Journal of the American Medical Association* 302(1): 67–72.
Sebald, W. G., (2002), *The Rings of Saturn* (London: Vintage).
Seneca, (2010), *Selected Letters* (Oxford: Oxford University Press).
Sessa, B., (2007), 'Is there a Case for MDMA-Assisted Psychotherapy in the UK?'. *Journal of Psychopharmacology* 21(2): 220–4.
Shakespeare, W., (1980), *Hamlet* (London: Penguin).
Shapsay, S. and Pimple, K., (2007), 'Participation in Biomedical Research is an Imperfect Moral Duty: a Response to John Harris'. *Journal of Medical Ethics* 33(7): 414–17.

Shook, J., (2012), 'Neuroethics and the Possible Types of Moral Enhancement'. *AJoB Neuroscience* 3(4): 3–14.
Singer, P., (1972), 'Famine, Affluence and Morality'. *Philosophy and Public Affairs* 1(3): 229–43.
— (1996), *Practical Ethics* (Cambridge: Cambridge University Press).
Smith, A., (1981), *An Inquiry into the Nature and Causes of the Wealth of Nations* (Indianapolis: Liberty Fund).
Stapledon, O., (2004), *Last and First Men* (London: Gollancz).
Stjernschantz Forsberg, et al., (2013), 'Why Participating in (Certain) Scientific Research is a Moral Duty'. *Journal of Medical Ethics,* doi:10.1136/medethics-2012-100859.
St John, C. and Heald-Moore, T., (1995), 'Fear of Black Strangers'. *Social Science Research* 24(3): 262–80.
Takala, T., (2009), 'Human before Sex? Ectogenesis as a way to Equality', in Simonstein, F., (ed.), *Reprogen Ethics and the Future of Gender* (Dordrecht: Springer), pp. 187–96.
Taylor, T., (2010), *The Artificial Ape: How Technology Changed the Course of Human Evolution* (New York: Palgrave Macmillan).
Thiele, L. P., (1994), 'Heidegger on Freedom: Political not Metaphysical'. *The American Political Science Review* 88(2): 278–91.
Timmerman, J., (2006), 'Kantian Duties to the Self, Explained and Defended'. *Philosophy* 81(3): 505–30.
Vedder, A. and Klaming, L., (2010), 'Human Enhancement for the Common Good – Using Neurotechnologies to Improve Eyewitness Memory'. *AJOB Neuroscience* 1(3): 22–33.
Voigt, C. and Keasling, J., (2005), 'Programming Cellular Function'. *Nature Chemical Biology* 1(6): 304–7.
Wagenaar, W., (2008), 'Enhancing Memory in the Criminal Trial Process', in Zonneveld, L., Dijstelbloem, H., and Ringoir, D., (eds.) (eds), *Reshaping the Human Condition: Exploring Human Enhancement* (The Hague: Rathenau Institute, in collaboration with the British Embassy, Science and Innovation Network and the Parliamentary Office of Science & Technology).
Wasserman, D., (2013), 'Devoured by our own Children: the Possibility and Peril of Moral Status Enhancement'. *Journal of Medical Ethics* 39(2): 78–9.
Wilkinson, S., (2010), *Choosing Tomorrow's Children* (Oxford: Clarendon).
Williams, B., (1973), 'A Critique of Utilitarianism', in J. J. C. Smart and B. Williams (eds), *Utilitarianism: For and Against* (London: Cambridge University Press).
— (1999a), *Problems of the Self* (Cambridge: Cambridge University Press).
— (1999b), *Moral Luck* (Cambridge: Cambridge University Press).
Wittgenstein, L., (2001), *Philosophical Investigations* (Oxford: Blackwell).
World Health Organisation, (2005), *Make Every Mother and Child Count* (Geneva: World Health Organisation).
Young, L., et al., (2010), 'Damage to Ventromedial Prefrontal Cortex Impairs Judgment of Harmful Intent'. *Neuron* 65(6): 845–51.
Zhang, F., et al., (2010), 'Fossilized Melanosomes and the Colour of Cretaceous Dinosaurs and Birds'. *Nature* 463(7284): 1075–8.

# Index