

Moral Bioenhancement and the Utilitarian Catastrophe

NICHOLAS AGAR

Abstract: This article challenges recent calls for moral bioenhancement—the use of biomedical means, including pharmacological and genetic methods, to increase the moral value of our actions or characters. It responds to those who take a practical interest in moral bioenhancement. I argue that moral bioenhancement is unlikely to be a good response to the extinction threats of climate change and weapons of mass destruction. Rather than alleviating those problems, it is likely to aggravate them. We should expect biomedical means to generate piecemeal enhancements of human morality. These predictably strengthen some contributors to moral judgment while leaving others comparatively unaffected. This unbalanced enhancement differs from the manner of improvement that typically results from sustained reflection. It is likely to make its subjects worse rather than better at moral reasoning.

Keywords: bioenhancement; utilitarianism; morality; moral reasoning

This article challenges recent calls for moral bioenhancement—the use of biomedical means, including pharmacological and genetic methods, to increase the moral value of our actions or characters. It responds to those who take a practical interest in moral bioenhancement. Among those with a practical interest are Ingmar Persson and Julian Savulescu.¹ They argue that moral bioenhancement is required to prevent ultimate harm, an event that they view as “making worthwhile life *forever* impossible on this planet.”² In Persson and Savulescu’s words, “Modern scientific technology provides us with many means that could cause our downfall. If we are to avoid causing catastrophe by misguided employment of these means, we need to be morally motivated to a higher degree.”³

I argue that moral bioenhancement is unlikely to be a good response to the extinction threats of climate change and weapons of mass destruction. Rather than alleviating those problems, it is likely to aggravate them. We should expect biomedical means to generate piecemeal enhancements of human morality. These predictably strengthen some contributors to moral judgment while leaving others comparatively unaffected. This unbalanced enhancement differs from the manner of improvement that typically results from sustained reflection. It is likely to make its subjects worse rather than better at moral reasoning.

The prospect of using biomedical means to alter inputs into moral thinking should be especially alarming to utilitarians. It threatens a long-standing compromise between moral common sense and the principle of utility. The possibility of moral bioenhancement seems to invalidate a popular response to the suggestion that utilitarianism demands that we perform repugnant acts. Prominent in such objections are cases in which utilitarianism would impose significant sacrifices on the few to benefit the many. One familiar scenario involves the forcible harvesting of organs from healthy people to ensure the survival of greater numbers of people who require transplants. A standard utilitarian response is to claim that we cannot

I am grateful to Felice Marshall, Vojin Rakić, and an anonymous referee for comments on this article.

be morally required to do what is impossible for us. Few human surgeons could bring themselves to forcibly harvest the vital organs of healthy innocents. Facts about the human psychologies of the surgeons give evidence that their attempts to do so are likely to cause the healthy donors to suffer but are unlikely to produce any of the planned benefits. It follows that human emotional and psychological limits spare utilitarians from having to endorse abhorrent conclusions. Moral bioenhancement offers the means to overcome these psychological limits. We, as we are, could not impose the repugnant sacrifices, but moral bioenhancement may permit us to alter ourselves so as to more easily perform them. Utilitarianism thus fails to evade the objection that it requires repugnant acts. I propose that we locate the problem not with utilitarianism but instead with the program of moral bioenhancement. We should acknowledge utilitarianism as a good, specifically human morality—a philosophically defensible theory of right and wrong tailored to human psychological limits. We should resist the suggestion that it become a template for the remaking of human moral psychology.

In-Principle versus In-Practice Objections against Moral Bioenhancement

It is important to distinguish in-principle from in-practice objections against moral bioenhancement. John Harris offers an example of the former kind of objection.⁴ He argues that moral enhancement via biomedical means undermines autonomy and the value of moral agency. The value of our free acts assumes a “freedom to fall” that comes with having the choice to act immorally. There is a conflict between moral deliberation and moral bioenhancement. According to Harris, the use of biomedical means to morally enhance directly affects our behavior without being subject to rational review. We lose the value associated with morally good actions.

In-principle objections such as that offered by Harris are philosophy’s stock-in-trade. It is therefore not surprising that Harris’s argument has prompted a spirited debate. One can reply to Harris’s in-principle objection by describing a case in which morally bioenhanced humans act in ways that are both better and free. If successful, such an argument would show that Harris is mistaken in finding an incompatibility between valuable free acts and moral bioenhancement.

The evaluation of objections in practice takes philosophers out of areas that are their exclusive intellectual preserve. It tends to require an understanding of the effects of our actions on the world. Someone who mounts an in-practice objection against moral bioenhancement can allow that it may be acceptable in principle. The in-practice objector allows that when one seeks to imagine moral bioenhancement leading us to act in morally better ways, one does not seek to imagine an impossible scenario. Some ways the world might have been or could be would lead the objector to endorse the practice. But the in-practice objector alleges that the world differs from these ideal states in ways that predictably lead to bad outcomes.

Defenders of moral bioenhancement should take care to respond appropriately to in-principle and in-practice objections. It is a mistake to seek to reply to an in-practice objection as if it were an in-principle objection. It is dangerous to bungee jump without properly strapping yourself into your harness. This is so even if successful bungee jumping without being strapped in is possible in principle. There is no law of logic that necessitates a bad outcome from an unrestrained bungee jump.⁵ The fact that there is no decisive in-principle objection against

unstrapped-in bungee jumping does not address the suggestion that it is dangerous.

In-practice objections emerge from a proper understanding of the context in which a given action is to be performed. Often this understanding will draw on the relevant science. A basic understanding of the properties of the human body and the effects of gravitational forces helps us to appreciate the danger of bungee jumping while not strapped into a harness. There is a high likelihood that the jumper will not maintain her grip and, as a consequence, will hit the ground with lethal force.

The evidence for this article's in-practice objection against moral bioenhancement comes from an account of how we make moral decisions. I do not suggest that an understanding of moral psychology directs that moral bioenhancement is necessarily bad. No law of logic directs the necessity of the ill effects I mention. But an understanding of how humans make moral decisions and the likely effects of biomedical alterations of these should nevertheless show that ill effects are probable. An understanding of human moral psychology suggests that there is an unacceptable risk of morally bad solutions to extinction threats.

The Moral Significance of Normal Human Psychology

A significant motivation for moral bioenhancement comes from the recognition that much of the reason for the mess we currently find ourselves in comes from our moral failings. We should allow that moral improvements could produce more effective action in respect to the climate crisis. Take your favorite environmental campaigner. It is not implausible to suppose that alternations that made the relevant parts of our moral psychologies similar to the psychologies of the environmental activists Rachel Carson and David Suzuki would lead to better behavior in respect to the environment.

To understand the danger of using biomedical means to intervene in moral reasoning, we need some understanding of how humans make moral decisions. The following account is necessarily schematic, but it should suffice for the purposes of this article. Human moral judgment draws on a disparate variety of mental abilities. It is informed by reason. Humans reason about the soundness of moral principles and the effects actions have on the world. Our affective capacities are also relevant. Emotional bonds seem relevant to both how we treat others and how we *should* treat them. The close emotional bonds between friends seem to be the basis of some mutual moral entitlements. Our behavioral capacities are relevant too. People are not morally required to perform actions that are impossible for them. It is plausible that there is a reduced obligation to perform actions that our awareness of our behavioral capacities informs us to be very difficult.

Exactly how this disparate collection of cognitive, affective, and behavioral inputs produces good moral judgment is the topic of intense debate among philosophers. There is a range of views about the relevance of each. All plausible accounts of moral judgment give each of these capacities some significant role.

A feature of moral bioenhancement is that it targets specific inputs into moral judgment. Some defenders of moral bioenhancement describe pharmacological agents that boost empathy. Oxytocin seems to have this effect, and this has led to some interest in it as a possible moral bioenhancer. Advocates of moral bioenhancement allow that these are early days for their project. A well-funded research

program should reveal compounds that boost a variety of contributors to moral thinking. We might look forward to a future containing biomedicines that increase the human capacity to empathize with those who suffer. There may be biomedicines that increase our capacity to avoid errors in moral reasoning. There may be biomedicines that increase human behavioral capacities.

One feature of such interventions makes them dangerous. They are *piecemeal interventions*. They target specific psychological influences on moral judgment. This piecemeal approach differs from the means by which moral improvement typically occurs. When successful, moral education typically strengthens many contributors to moral thinking. A change in your affective responses is typically accompanied by changes in the moral reasoning you perform. It is the piecemeal nature of changes to moral inputs produced by moral bioenhancement that makes it dangerous.

How might the piecemeal approach to boosting inputs into moral judgment pose a danger to moral thinking? I begin with a methodological point. Normal human psychology plays an essential role in justifying our moral principles. It gives rise to the intuitions against which we compare candidate moral judgments.

If we suppose that normal human psychology plays a justificatory role in morality, then it is clear how deficits in these capacities tend to lead to bad moral behavior at the same time as impairing moral judgment. They worsen our ability to grasp the criteria for correct moral judgment. Piecemeal interventions are more likely to be successful when their aim is to bring individuals with cognitive, affective, or behavioral deficits up to levels properly considered normal for human beings. Here interventions function much like injections of insulin for diabetes—their purpose is to restore biological normalcy. They are likely to be much less successful when directed at individuals who achieve moral normalcy. Because the climate crisis results from the individual and collective errors of people who achieve moral normalcy, moral bioenhancement is unlikely to be of any use.

What is the concept of moral normalcy at work here? Morally normal people have the cognitive, emotional, and behavioral capacities required to understand and comply with moral judgments. In an earlier presentation of this claim, I proposed that the infamous Soviet dictator Josef Stalin might be morally normal in this sense. This suggestion seemed to scandalize Persson and Savulescu.⁶ To say that Stalin could have been morally normal is not to praise him. Indeed, it should be part of a moral condemnation of him. If he was morally normal, he was capable of understanding the wrongness of his actions. He does not get the excuse that applies to crazed murderers who inflict suffering but lack any insight into the moral wrongness of their actions. Moral normalcy is, as Persson and Savulescu point out, “far from good enough.” But that is precisely the point. A morally normal person has the capacity to act morally. There’s a further issue as to whether she does so. The fact that moral normalcy does not suffice to ensure morally good behavior should not be taken to suggest that it is not an important practical precondition for morally good behavior.

Consider some consequences of failing to achieve moral normalcy. Suppose that philosophers sought to test their principles on people with extreme empathy deficits. Suppose also that empathy deficits are not accompanied by any impairment in moral cognition. Someone with normal powers of moral cognition but with a very low capacity to empathize may be fully convinced by John Stuart Mill’s arguments for the principle of utility. Such a person may be less troubled by the philosophical

thought experiments appealed to by utilitarianism's opponents—thought experiments in which directly inflicting suffering leads to less suffering overall. When people with significant empathy deficits imaginatively place themselves into such situations, they might find it easy to imagine doing as the principle of utility directs. This seems to describe the psychology of the fictional TV psychopath Dexter Morgan. His lack of empathy enables him to commit utility-maximizing murders impossible for emotionally normal human beings. Dexter's job in the police department enables him to arrange the deaths of murderers who are likely to murder again. We can imagine that if Dexter were presented with a thought experiment in which he was asked to surgically remove the vital organs of a single healthy person to save the lives of five others, he might rather easily be able to imagine himself doing this and therefore accept the action's moral correctness. We don't need to imagine that Dexter completely lacks empathy. A utilitarian whose feelings of empathy are genuine but less strong than those of normal humans may find it easier to comply with what he takes to be the correct moral theory when it conflicts with relatively weak moral feelings.

Consider someone with subnormal levels of moral cognition. This individual is incapable of recognizing what to most of us would be obvious moral consequences of his actions. We need not suppose that an individual who has a deficit in moral cognition is ignorant about key aspects of the physical or social world. For example, a person with subnormal levels of moral cognition may know basic facts about why people receive benefit checks. He may also know that those who need these payments but don't receive them suffer. But the person with impaired moral cognition may fail to register these facts as morally relevant. He may say, "So what if those I steal from suffer." He is thus easily able to self-justify stealing the checks when doing so enables him to achieve an emotionally salient moral end, such as benefiting dependent children. A less extreme version allows that he believes that it is, in general, morally good that benefit checks go to people in need. But this act of moral cognition exercises such a weak motivational influence that he finds it easy to override it when doing so benefits his children.

These deficits are likely to have an effect not only on the actions that we perform but also on actions that we judge to be morally justified. If philosophers presented moral thought experiments exclusively to collections of individuals whose moral emotions were deficient or to collections of individuals whose moral cognition was subnormal, then they are likely to arrive at principles unpalatable to most of us. People with normal moral cognition, emotions, and motivational capacities are the reference points for moral claims that philosophers accept as justified.

Morally normal humans judge that it is morally permissible to place their children's welfare ahead of the welfare of a slightly greater number of strangers. But they judge that it would be morally wrong to deliberately torture greater numbers of strangers so as to somewhat enhance the welfare of their children. These and other intuitive moral judgments draw on the cognitive, affective, and behavioral capacities of normal human beings. Human parents have special moral responsibilities to their children partly in virtue of strong ties of emotion. Yet we understand that we have obligations to strangers. There is a particular pattern of trade-offs between moral cognition and moral emotion that produces acts that we judge to be morally correct.

The intuitions of morally normal humans are key here. The moral judgments that normal human beings recognize as intuitive are unlikely to be so recognized

by someone severely deficient in empathy or by people whose cognitive impairments prevent them from understanding that suffering experienced by a stranger might be just as intense as suffering experienced by a friend. Piecemeal enhancement of inputs into moral judgment can have a similar effect. It is likely to produce judgments that depart from the compromise between affect and cognition endorsed by moral common sense.

Suppose that a piecemeal moral enhancement boosts empathy. Its strengthening of emotional bonds is likely to shift the point of compromise between the deliverances of moral reason and the deliverances of moral emotions. Suppose your moral reasoning directs you to a consequentialism that finds no moral difference between the interests of your children and the interests of strangers. You nevertheless experience emotional bonds with your children that you do not experience with strangers. Your considered moral judgments are compromises between rationally persuasive principles and feelings of attachment. If piecemeal moral enhancement boosts empathy, then you are more likely to judge that benefits to your children justify imposing sacrifices on strangers. Remember, this is an in-practice objection against moral enhancement. The fact that it is, in principle, possible that humans with heightened empathy will not judge that it is morally permissible to impose sacrifices on strangers does not speak to the concern that they are more likely to do so.

Suppose that the piecemeal enhancement of moral cognition increases your capacity to comply with a moral principle that you rationally endorse. You are likely to be better at resisting the moral emotions that speak against this. You may come closer to imposing sacrifices on your children when the principle of utility directs that you should. A strengthened capacity to comply with moral principles that you judge to be reasonable may lead you to overcome affective restraints. If inflicting harm on your children promotes overall happiness, then you may be more likely to do it. I will have more to say about destabilizing implications of piecemeal moral bioenhancement for utilitarianism in the next section.

When directed at people who achieve moral normalcy, piecemeal moral bioenhancement is likely to lead to excesses that impair moral judgment. Again, remember that this is an in-practice objection. There is no reason to think it entirely impossible that a piecemeal moral bioenhancement or a combination of piecemeal moral bioenhancements might succeed in producing exactly the kinds of moral improvements endorsed by widely held moral intuitions. There is no reason to think it impossible to survive bungee jumping without a harness. Both acts are, nevertheless, ill advised.

One reason that balanced moral bioenhancement is difficult to produce by artificial means is that the inputs into good moral judgment are so different. Suppose you seek to improve your marathon times by enhancing the strength of your lower limbs. You know that the enhancements should be balanced. A strengthening of your right limb that is unmatched by a strengthening of your left limb will worsen rather than improve your marathon performances. Here it is clear what counts as balanced enhancement. You ought to produce improvements of the left limb that are symmetrical with those produced in the right limb. Symmetry is no guide for those seeking balanced moral enhancement. The inputs into moral judgment are very diverse. The superior judgments of those we acknowledge as our moral betters result from achieving a complex relationship between affective, cognitive, and behavioral inputs. We have some sense of what this enhanced balancing would be

like from the inside but little idea about how to achieve it by artificial means, from the outside.

The suggestion that moral bioenhancement tends to be piecemeal, and therefore dangerous, does not lead to the conclusion that we could be content with the moral status quo. Moral education and reflection were instrumental in many moral successes, including the abolition of slavery. The freeing of the slaves did not occur through the piecemeal enhancement of selected inputs into moral judgment.

The Catastrophic Implications of Moral Bioenhancement for Utilitarianism

Utilitarianism is often presented as too demanding a morality. According to this objection, complying with utilitarian morality prevents one from leading a worthwhile human existence. The principle of “ought implies can” is frequently called on to excuse utilitarians from having to meet apparent extreme demands. Utilitarianism supports no moral requirement to perform impossible acts. Some apparently good acts that seem to fall within the scope of human physical capabilities turn out to be impossible because of psychological and emotional limits. This has become a stock utilitarian response to the theory’s abhorrent apparent implications.⁷ Drugs that alter human moral dispositions are bad news for this compromise between the principle of utility and moral common sense.

Consider the following scenario. Angela is a loving mother of a child and is also a medical researcher seeking a better treatment for HIV. According to the World Health Organization, approximately 1.7 million people died of AIDS-related illnesses in 2011. Angela conceives of a new therapy to treat the disease—therapy X. Therapy X implements a radically new idea about how HIV/AIDS should be treated. If this idea is vindicated, we should expect to see dramatic improvements in the treatment of HIV+ people. This is not certain. Angela understands that many experimental therapies fail to produce their hypothesized benefits. The history of medical research contains many in-principle cures for serious diseases that fail to lead to any effective treatments. Therapy X must be tested. Angela is aware that delays in testing can cost lives. Delays can be morally costly. Approximately 5,000 people die of AIDS-related diseases every day. The testing of new therapies is often beset by bureaucratic delays. Angela understands that, should therapy X realize her hopes, then it will save many lives and that each day of delay costs thousands of lives. Should Angela bring therapy X home and test it on her child? Testing it will require that she first infect the child with HIV.

We must stipulate a few aspects of the story to make it a true test of the possible effects of moral bioenhancement on utilitarian motivations. There are many people already infected with HIV who might be more suitable participants in a clinical trial of therapy X. However, submitting the drug to the standard approval process will take time. Every day of delay costs thousands of lives. Suppose that X either cures HIV or significantly reduces the severity of its symptoms and the likelihood of its transmission from one human to another. Suppose that sending the drug through the proper channels delays its availability by 10 days. The cost of that delay could be many tens of thousands of lives. Angela’s plan to test the drug on her child plausibly accelerates its arrival. If the experiment yields the hoped-for results, Angela can work out how to forcefully present to the authorities her strong hunch that therapy X really will be a much-hoped-for breakthrough in the

treatment of HIV/AIDS. She understands that the experiment will harm her child. But she also understands that it should bring closer a therapy that will relieve the suffering of millions. She feels the force of utilitarian reasoning supporting her illegal experiment.

Note that therapy X does not have to either be an effective therapy or directly lead to one for the experiment on her child to predictably relieve a large amount of suffering. Suppose therapy X is a dead end in the search for an effective treatment for HIV/AIDS. Angela's experiment should enable the initially promising therapy to be removed from consideration, allowing attention to quickly shift to better avenues of research. Angela is a very talented medical researcher. The gain here should be measured in terms of the time we would expect conventional methods of medical research to demonstrate the ineffectiveness of therapy X when compared with the time it will take Angela to arrive at this conclusion if she tests it on her child. The net result should be the speedier arrival of enhanced treatments for a disease with an annual death toll of more than one and a half million lives.

On the face of it, the principle of utility would seem to demand the test. The high probability of harm suffered by her child must be appropriately balanced against a low probability of benefit for millions. Decision theory tells us that it can be right to incur probable small losses so as to achieve less probable benefits of greater magnitude.

Fortunately for utilitarians, the principle that ought implies can absolves them from having to endorse the experiments. Suppose that Angela has a normal human combination of moral emotions, reasoning, and capacities. Angela will predictably be unable to successfully conduct the experiment. She loves her child. Any attempt to use her child as an experimental guinea pig is likely to have the effect of inflicting suffering without producing medical research of any value. No researcher with normal human psychology could successfully enhance understanding of HIV/AIDS by experimenting on her loved ones in this way. Therefore these experiments cannot be morally required. Indeed, Angela's normal human psychology means that utilitarianism requires her not to attempt to conduct the experiments. She will predictably inflict much suffering on one person without producing any benefit for others. Utilitarians evaluating her behavior can be confident that she has done the right thing in accepting the delays and following the conventional channels in testing therapy X. Extraterrestrial medical researchers might find themselves subject to a requirement to test therapy X on innocent healthy humans. But we are not. We have thus resolved an apparent tension between moral common sense and the principle of utility.

I suspect that the capacity to alter moral psychology in a piecemeal fashion invalidates this compromise between moral common sense and the principle of utility. Suppose Angela has access to a piecemeal moral bioenhancer that will predictably strengthen the grip on her psychology of the principle of utility—the moral principle that she finds most credible. In such circumstances, the principle that ought implies can offers her (and her child) reduced protection against utilitarian demands. Having self-administered the necessary moral bioenhancer, Angela commences the experiment confident that her feelings of attachment to her child will not interfere with her moral mission to reduce global suffering. Inflicting suffering on her child now becomes an effective means of maximizing happiness.

It is not absurd to suppose that a moral bioenhancer might work this way. An agent that enhances our ability to comply with moral principles that we endorse might work in a way that is similar to the way in which beta-blockers enhance the performances of classical musicians. Beta-blockers can help to steady a violinist's hand so that her nerves do not interfere with her public performances. We can imagine a biomedical moral enhancer that works in an analogous way. A utilitarian recognizes that the sacrifice of one of his nearest and dearest would maximize happiness. But his normal human psychology prevents him from being able to successfully impose this utility-maximizing sacrifice. He may achieve the proximal end of causing suffering to his nearest and dearest, but he is unlikely to be able to do so in a way that produces the distal end of good consequences. A talented violinist takes a beta-blocker to ensure that her nerves do not interfere with her public performances. She can perform at a level that better reflects her talent. A utilitarian takes a piecemeal moral enhancer that predictably suppresses emotional inhibitions preventing him from successfully maximizing happiness. The drug strengthens the grip on his psychology of a moral principle that he believes to be true. He is now better able to maximize happiness.

Utilitarians who presently possess normal human psychologies should be terrified by the wide vistas of possibilities opened up by ways to systematically alter human moral psychology. These piecemeal interventions have a potentially disastrous impact on a long-standing utilitarian compromise with commonsense morality. Piecemeal interventions in human moral dispositions permit actions impossible for unmodified humans.

Suppose Angela is a utilitarian who has access to a piecemeal moral bioenhancer that will predictably enhance her capacity to produce good consequences. Should she take the drug, thus enabling her to experiment on her child with therapy X? The outcome in which she takes the moral beta-blocker and successfully performs the experiment seems preferable in utilitarian terms to the outcome in which she does not take the drug and is therefore unable to perform the experiment. The misery she may feel at having caused her child to suffer must be balanced against the reductions in suffering that result from an accelerated progress in the search for better therapies for AIDS.

Utilitarianism as a Distinctively Human Morality

I think that there is a way of understanding utilitarianism that does not lead to the conclusion that Angela should take the moral beta-blocker and perform the experiment. We should acknowledge utilitarianism as a distinctively human morality. It gives moral advice to individuals with normal human psychologies. It says nothing about how beings who are emotionally and psychologically very different from us should act. It says nothing about how we should modify our moral psychologies.

We can compare utilitarianism as a human morality with exercise programs designed for human beings. Most of us could have greater levels of physical fitness than we currently do. The purpose of an exercise program is to enable us to improve our levels of physical fitness. Good exercise programs are tailored to the limits of human physiology. A program that recommends that you spend your day eating pizza and watching zombie movies is a bad one because it does not achieve the end of improving your physical fitness. An exercise program that

recommends that you spend your first week running daily ultramarathons is bad for a different reason. It is improperly informed by the limits of human physiology. There are possible beings for whom such a program would be effective at improving physical fitness, but those beings are not us. There is a question about whether, if the means existed, we should become such beings. It seems to me that this question is beyond the scope of a trainer who conceives of his job as helping human beings to become fitter. He can legitimately reply that he has nothing to say in response to this question. His expertise is in the physical fitness of human beings. He can recognize the enhanced beings as, in some sense, superior in physical prowess to humans while rejecting the suggestion that criteria relevant to them have any relevance to the recommendations he makes.

An analogous reply is available to a utilitarian who conceives herself as having the task of giving moral advice to human beings. She can allow that utilitarians unimpeded by the constraints of human psychology will perform actions that are in some relevant respect morally superior to our own. But suppose she views herself as an expert in moral theories tailored to the limits of human psychology. She might accept utilitarianism as a good human morality while rejecting suggestions that the theory become a template for the redesign of our moral psychologies. There is no inconsistency in recommending utilitarianism as a morality for human beings while rejecting the suggestion that human psychology be systematically modified so as to permit outcomes that seem superior in utilitarian terms. This is not to say that utilitarians should be completely indifferent to the states of moral psychologies. They can allow that forms of empathy training that lead to improved responsiveness to human suffering are good in utilitarian terms. But this awareness of the effects of moral psychological improvement can and should be informed by human limits. It need not extend to the endorsement of artificial means.

The Proper Philosophical Place for Speculations about Moral Bioenhancement

This article has presented a practical response to moral bioenhancement. What we know about human moral psychology, and what we can predict about how moral bioenhancers will work, suggests that they should not be used to make morally better humans. The prospect of piecemeal moral bioenhancement should be especially worrying to utilitarians. Their mode of operation threatens a popular response to the claim that utilitarianism requires repugnant acts. The selective suppression of human emotions opens up the possibility of repugnant utility-maximizing acts. I suggested a way to respond to this possibility. Utilitarians can consistently reject moral bioenhancement if they view their ethical theory as a distinctively human morality. It gives good moral advice to humans while being silent on the question of how humans might alter their moral dispositions.

What does this mean for philosophical speculation about moral bioenhancement? Is there a way for philosophers to justify an interest in moral bioenhancement without recommending that we seek to practice it? I propose that moral bioenhancement scenarios belong in thought experiments whose purpose is to support or challenge philosophical claims about enhancement.⁸ In this way they resemble Robert Nozick's famous experience machine thought experiment. Nozick imagines a machine capable of delivering all manner of good experiences to those who are prepared to plug into it, substituting pleasurable artificial experiences for less pleasurable veridical ones. He treats our reluctance to plug in as indicating that

we strongly value a connection with reality. It stands to reason that Nozick should not be interpreted as advocating the construction of experience machines. But even philosophers who seek to reverse Nozick's conclusion can distinguish their theoretical interest in the experience machine thought experiment from a practical interest in the construction of experience machines. It is possible to view the thought experiment as demonstrating the truth of hedonism without going on to say that the truth of hedonism recommends a global program of constructing experience machines. Philosophers should take care to distinguish an interest in thought experiments about moral bioenhancement from a misguided campaign to bring moral bioenhancement to the people.

Notes

1. Persson I, Savulescu J. *Unfit for the Future? Modern Technology, Liberal Democracy and the Need for Moral Enhancement*. Oxford: Oxford University Press; 2012.
2. Persson, Savulescu 2012, at 46.
3. Persson, Savulescu 2012, at 8.
4. Harris J. Moral enhancement and freedom. *Bioethics* 2011;25:102–11.
5. Indeed, this clip demonstrates that survival is possible: Russian daredevil stunts. *9 News.com*; available at <http://news.ninemsn.com.au/world/2013/11/20/11/15/russian-daredevil-stunts> (last accessed 22 May 2014).
6. They say, "The moral normalcy that Agar appeals to (his example of it is, incredibly, Joseph Stalin) is far from good enough." See Persson I, Savulescu J. Reply to commentators on *Unfit for the Future*. *Journal of Medical Ethics* 2014 Jan 10; available at <http://jme.bmj.com/content/early/2014/01/10/medethics-2013-101796.full> (last accessed 5 Nov 2014).
7. For a prominent recent presentation of this response, see Greene J. *Moral Tribes: Emotion, Reason, and the Gap between Us and Them*. London: Penguin Press; 2013.
8. Indeed, this is how Tom Douglas uses them. See Douglas T. Moral enhancement. *Journal of Applied Philosophy* 2008;25:228–45.