Published in the Cambridge Quarterly of Healthcare Ethics

Volume 24 / Issue 01 / January 2015, pp 8-22

Special Section: How Moral Is (Moral) Enhancement?

A Debate about Moral Enhancement

JOHN HARRIS and JULIAN SAVULESCU

This is an edited transcript of a debate between John Harris and Julian Savulescu that took place before an audience during the "Enhancement: Cognitive, Moral and Mood" conference held at the University of Belgrade in 2013. The debate begins with the questioner asking John Harris to briefly define "enhancement."

Harris: We all possess many powers and capacities. I'm using one of many of those powers now. I'm using the power of speech. I can also see you. If you talk to me, I can hear you. These are powers that most human beings have. We can improve these in various ways. A woman sitting in the front row is improving her sight by wearing spectacles—that is an enhancement technology. It improves sight; so do microscopes. If you look down them, you get a magnification far beyond that available to a pair of spectacles. That, again, enhances your sight. So enhancement is a general name we use for improving upon all of the powers and capacities that we humans have or might potentially have. But this also includes other important powers and capacities like our ability to resist infection or resist disease, or indeed to augment those powers beyond the normal.

If it wasn't good for you, it wouldn't be an enhancement. So obviously one of the constraints, one of the things beyond which we will not go, we will not recommend things that will not be good for you. If you like an epigrammatic way of expressing this truth, there is a saying in English, "You can't have too much of a good thing!"—because if you have too much of it, it ceases to be a good thing. So you enjoy eating, but if you eat too much, you get too fat. It ceases to be a good thing. It doesn't mean eating is bad. It means that you can do it to excess. So whatever the improvements we try to make to powers and capacities, when those improvements cease to be good for us, then they cease to be enhancements. So we're talking about those improvements that we would like to have and that would make our lives better and more successful.

Questioner: Professor Savulescu, a question for you. What is moral enhancement?

Savulescu: As John said, we can enhance many aspects of our lives. This includes enhancing our relationships, including our marital relationships, but we can also enhance our relationships with others, including our moral behavior as human animals. One of the most important behaviors that we share with other mammals is ways of organizing ourselves and cooperating with each other and not harming members of our group. All human beings have these capacities but to varying degrees. We're living in a very exciting time, probably the most exciting and also dangerous time, because our powers through technology are ever increasing. Not only have we split the atom and created nuclear power, we've sequenced

the human genome, and human cloning is all but possible. But we are also starting to understand the science of morality, what makes people different, why some people commit crime and others don't, why some people are more empathetic or sympathetic, and others are less, and why some people are altruistic and others are selfish. This is not to say moral behavior is in only our genes or in our nature. It's to say that our nature can influence this just as the world in which we live influences us. Science is giving us also the capacity to change this biological part, this internal part, just as we can change the way we think and feel. We can also change how we morally behave towards each other, and this I think is a very exciting prospect that I think there are very strong reasons to research further.

Harris: Nature is morally indifferent. It can be good. It can be bad. What's natural is neither necessarily good nor bad. We would not practice medicine if we didn't believe in attempting to frustrate the course of nature. People naturally fall ill. They are naturally attacked by viruses or by bacteria. If we only did what was natural, we would not practice medicine. We would not do science. We would not do a whole range of things. So there is no magic in what's natural or unnatural. What we have to do is decide between those things which are natural, if you can define those, and those that are unnatural. We need to decide which are good for us and which are bad for us and which are useful and which are threatening. That is the only issue in my view. I don't disagree with Julian about that.

Savulescu: In many cases, people think that natural influences act mysteriously or in a different way to man-made enhancers, but they change your biology. For example, you can perform better in sport by training at altitude because you reduce the amount of oxygen in the air and your body produces more red blood cells. And you have a competitive advantage. Now, that's natural; it's completely legal in sport. You can do the same thing by donating some of your blood now and, in three months' time, when your body has replenished that, having that blood transfused back. Both of them increase the number of red cells. One is natural, the other is said to be unnatural, but both of them have exactly the same effects, and to think that natural influences are somehow different from unnatural ones is a deep mistake—we are a part of nature, and what we construct is a part of nature. When you learn something, a skill like riding a bike or playing the piano, your brain changes, and it changes irreversibly, just as if you gave the person a drug or changed the structure of the brain intentionally and directly. So I think it's this distinction between natural and artificial is very deep in people's psychology because, as animals, we liked to have green around us. We like to have abundant nature for our survival, so we are programmed to have this positive response to nature. Sorry, but I think we need to go further than that now and ask other questions.

Harris: Just to give you one more example, Julian has given some very good examples, but one of my favorites is something that we're all benefiting from at this moment. It's completely unnatural. It's synthetic. It's synthetic sunshine. Before we had synthetic sunshine, we humans used to go to sleep when it got dark and could get up and work when the sun came up, but with the advent of synthetic sunshine—firelight, candlelight, electric light—our lives changed. We could work and play through the hours of darkness. In this room there are no windows. We are bathed in synthetic sunshine, but I hope nobody thinks that this is a disaster.

Savulescu: The whole of human history has been about trying to change the world around nature to make our lives better: agriculture, farming animals, building dams, building shelters, clearing land, and so on. Now we have the ability to change not just the external world but ourselves to have better lives. This is done, as John said, through medicine, but now we can change other things that aren't diseases but that affect how well our lives will go, and that I think raises very important opportunities but also difficult questions.

Questioner: My question is for Professor Savulescu. We all know scientific change is occurring exponentially, but has our morality managed to keep up? If our morality has not kept up, should moral enhancement be compulsory?

Savulescu: If you're a scientist and you study the human animal (not some idealization or romantic version) and its behavior, you will find that we are essentially the same sort of animal we were two hundred thousand years ago. Our bodies haven't changed in that time, and indeed, our minds have not changed. We have the same sorts of dispositions, the same sorts of character, as our original hunter-gatherer ancestors, but in the last ten thousand years our society has radically changed, with farming and the ability to domesticate animals and create complex cities. In the last one hundred to two hundred years, we've had the industrial revolution and then the information revolution, the genetic revolution, the nanotechnology revolution. We've had so many revolutions in technology that the world is completely different. We still have the sets of dispositions, but our world is completely different. So what is it like to be a human being anywhere in the world whether you're Serbian or American? You have a basic set of moral dispositions, what's called moral psychology, a set of ways in which you behave towards other people. Our moral psychology is characterized by a disposition towards violence, especially towards other people outside of your group. You care about people, but mostly your family and friends and people in your small community, your group of one hundred and fifty. You saw the consequences of this kind of concern for people who are like you in your group when Yugoslavia broke down. We can cooperate, but in a limited way with people when they can see us and at one time. But the world today gives people not just automatic weapons. It gives them biological weapons. Within a decade, you'll be able to create the most powerful biological weapons, like smallpox, in a backyard laboratory. Now it only takes one out of seven billion people to decide, instead of shooting one hundred people with an automatic rifle, that they will do this. So we have a power that we've never had before, and we don't have the moral qualities to deal with that sort of power. We also won't be able to deal with things like climate change or global poverty and inequality, because we're not set up psychologically to deal with these sorts of problems. We don't think about the distant future. We don't think about people in a faraway place. We can't empathize with a large number of people, only with single individuals. That's in our nature. So we're not morally set up for the challenges of collective action problems like climate change or global poverty or with the threat of mass violence. So because of this and our understanding of the ways in which you could influence those, not just by surveillance or more police force, but by changing people, I believe that we have a strong imperative to investigate how we can use science not just to make people live longer, be smarter, but also be less prone to this kind of mass violence, more altruistic, more empathetic, more cooperative, with a greater sense of justice.

Should moral enhancement be compulsory? The easiest thing is to make any new technology or power voluntary, because if you let people choose for themselves, and they choose it, it's because they think it's better for themselves or better that they should do it. And that, I think, is a very useful strategy. This is done when pedophiles are offered chemical castration, which is a crude form of moral enhancement to reduce their sex drive as a way of getting earlier parole. They make the choice. They can stay in jail, or they can have this moral enhancement. I think that is something most people would think is an option that should be offered to people. But if the intervention is very effective and safe, and uncontroversially good, we should do it compulsorily. Education is an example. Education is given to everyone because it is an enhancement that is clearly good for people and good for society. It's also a form of moral education where values are instilled in children. Now transcranial electrical stimulation is one form of enhancing numeracy and other learning skills. It involves a light electrical current over the surface of the skull. It could also help you to learn the sorts of morals skills and qualities that we want in our children. If it enhances that, it may be a real question whether this is treated just as education is, and indeed as a part of education.

Questioner: Professor Harris, you're skeptical about the value and methodologies of moral enhancement. You said, my own belief is that the most reliable and effective method of moral enhancement both for now and for the foreseeable future is in fact an enhancement of our cognitive and intellectual powers.

Harris: I did say that. So far, I agree with almost everything that my colleague Julian Savulescu has said. I agree about the dangers we face. I agree that we need to bring a moral consciousness as well as our physiology up to date with the increasing power that we are able to exercise over the world and over one another. So far we're in agreement. I also agree with him that if we can do it safely and effectively, then there may be good reasons to make it compulsory, and I agree with the analogy that he used with education. Your country, I imagine, like mine, has compulsory education up to a certain age. We do that for a number of reasons. One, because we want our children to benefit from education, to understand the world better, to be able to function better—to indeed be enhanced by the education that they received; but we also do it to protect them and to enable them to provide for their futures and, increasingly, this is important to me, for my future. I'm increasingly dependent on the young for lots of things that I need. So far we are in agreement. The issue is over safety and efficacy. At the moment, the sorts of interventions that are available to affect things relevant to our moral nature are very simple attitudes that are sometimes called "prosocial attitudes"—the sort of attitudes that enabled our very distant ancestors to start cooperating with one another and, in a rudimentary way, caring for one another. Our prosocial attitudes are the ones that make us want to help our fellow creatures, usually our own species, often much more local than that, our own families first and possibly then those in our own village and then possibly those in our own society. Now, that's great, but I don't think it is actually the chief need of the contemporary world. We now live in a global village. It is not enough to improve the sympathy, the empathy, the caring behavior that we offer to our friends and neighbors, to our families and our workplace. We need more imagination, more awareness of how to generalize those very important feelings of sympathy, empathy, and cooperation. Not only across whole nations, states, and continents but right across the world, because our safety depends on

everybody's safety. Now, I am concerned that the sorts of interventions that we currently have available are likely to increase my willingness, if somebody were to fall down in front of me in pain, to rush and help them, but they might decrease my awareness that, although you seem to be in some pain, it's not very serious and there may be people in greater need who I can't see. And if I'm thinking about how to deploy those extra resources that societies have to deploy for caring for others, I might be better advised to deploy them elsewhere in Serbia where we now are, or elsewhere in Europe, or elsewhere in the world. So what worries me is that we will increase our capacity and our willingness to help those near to us at the expense of those further away, and it seems to me that what we need in the global village that we now all inhabit is the power and the willingness to generalize our affections and concern for others right across the planet, right across the globe. And so I think we are beginning to reveal not a difference in about what matters, but a difference in whether or not the ways that are currently available to improve elements of what we might think of as our morality are likely to work. Take a different thing; we've talked about sympathy and empathy. Take something like violence. It's generally felt that we are too violent a species (and that may be right), and other drugs are available which decrease our propensity for violence, which sounds like a good thing, but sometimes we need violence to defend ourselves and to defend those that we care about.

Questioner: So, for example, if a mother kills a person trying to attack her child, in order to save her baby's life. This kind of murder would be justified.

Harris: We would all think that was justified, and I don't want to decrease your capacity as a mother, or potential mother, to protect your children by making you incapable of the level of anger and resentment against threats posed to them, a capacity for rage and violence that you might need to protect them. We need a mixed bag of dispositions. We need violence as well as lots of other things.

Savulescu: These issues that John is raising are unavoidable, but we have to face them. To take his example of empathy, the fact is that there is a decline in how empathetic people in the United States are, and I think it's probably the same in Serbia. Empathy in the last thirty years amongst American college students, anyway, has been declining, particularly in the last ten years. So you have to ask: is this a good thing or a bad thing, and is the current level that we have too high or too low? If it's the case that empathy is actually a bad thing—we have too much of it—we have to ask whether maybe we should reduce it. So we have to address these questions because (a) we all differ; (b) these sorts of properties change; (c) we're increasingly able to make changes to ourselves. And as John agreed, evolution and nature don't deliver up the right solutions to our problems. We can't say, "Oh let's just let nature sort it out," because it's very likely that nature hasn't selected the right distribution of, for example, psychopaths in the population. One percent of people are psychopaths. It is not that this is somehow ideal for the sort of world we're living in today. This is just an artifact of the way in which we evolved. So we have to ask: maybe we need more. Maybe we need less.

Harris: How many psychopaths do we need?

Savulescu: People have given this argument: Psychopaths can be leaders of big businesses. They are often the hardheaded individuals, but then when you look at the data, in fact they often ruin business. So you cannot answer this question—"Is it good for society to have psychopaths or not?"—without doing scientific research, just as you cannot look at what level of empathy is best going to deal with the problems that we agree we face. You need both science and ethics: ethics to set the values; science to tell you what will achieve them.

Questioner: I'm sure that we are all aware of the gap between what we believe to be right and how we act. How can we ensure that even if we undergo moral enhancement, we will act upon knowing what is right? What can bridge that gap?

Savulescu: We already do moral enhancement in lots of different ways, and, in fact, one of my Ph.D. students, Sylvia Terbeck, showed that the common drug propranolol, used to treat high blood pressure, reduces implicit or subconscious racism. But here's a very common example of a moral enhancement that is performed on children to increase their motivation: the use of drugs like Ritalin for attention deficit disorder. One of the worst things you can have as a human being is an inability to control your impulses and to delay your gratification until later, so you act impulsively. This leads you to be violent. It leads you to have great difficulty learning, making friends, and you end up at the very bottom of the economic scale. Now, attention deficit disorder is such a disorder of self-control. When Ritalin is given to children, it's in part to increase their ability to control these impulses and increase their motivational strength to do what they know is the right thing to do, and it's in part to stop those children being violent to other children as well as for their own good. So this is a very common drug used already that is partly a moral enhancer.

Harris: I'd hate to have too much agreement, because it's not nearly as much fun as disagreement, but again I agree with you, particularly about Ritalin. Ritalin is also a very powerful cognitive enhancer. It helps the same children who not only have problems with the focusing of their attention and with their restlessness and sorts of symptoms that Julian has described, but they also do poorly at school for those reasons. They can't concentrate. They can't organize themselves. They often find difficulty with their handwriting and a range of other skills as well. Now, Ritalin improves all of those things, and it improves their capacity to succeed at school, this grew up to about 10 percent with Ritalin. That's a big improvement. So again, I'm in favor of the use of Ritalin. For me, it's more important use it as a cognitive enhancer, because it enables kids to benefit from going to school much more. They succeed, they're happier as a result of that, so it has many very good effects. It also improves another range of capacities which we're going to need to face, the dangers that Julian has talked about and calls "ultimate dangers." Things like global warming, the risk of chemical and biological warfare, and the risk of new infectious diseases which are always arising in nature. Nature is probably the biggest bioterrorist that we face.

So we're going to need to improve our cognitive skills, and Ritalin is one of the drugs that may help us to do that; but this raises an issue that we do disagree about, and it's very well worth discussing. Julian has said, on a number of occasions, that because science and technology are so rapidly advancing and some things that science and technology are producing are dangers—the capacity, for example, to engineer viruses that haven't been communicable by aerosol, that is, in a way you can just breathe in, into viruses that are

aerosol transmissible. Two groups in the last eighteen months, one in America and one in the Netherlands in Holland, have engineered the H5N1 strain of bird flu so that it is communicable just by breathing the air that these birds have breathed, whereas formerly it was only communicable by contact with infected birds. The scientists did this because they realized that this form of flu might mutate in nature and present this danger. But the problem was, of course, that once you've learned how to engineer this change in the laboratory, then terrorists can do this in their back room and make a terribly powerful weapon for widespread terror. Now, Julian is worried, and I am too, that science is increasing the capacity for evil people to do harm by creating new diseases in their kitchens and spreading them, and he wants to put a brake on scientific development until moral enhancement, moral improvement, is sufficiently widespread for us to be confident that we can let loose the scientists in the laboratories again. This worries me greatly, because I look to science to solve many of the dangers that face us at the moment. I look to science not only to help us with new communicable diseases, but to produce better moral enhancers than we have at the moment. I don't think we can put the brakes on science until we have universal moral enhancement. I know you want to interrupt me. May I have one more sentence?

Questioner: [Laughs.] Okay, one more.

Harris: In the future there will be no more planet Earth, and there will be no more human beings; but we don't need to worry about either of those two things, because I believe when our sun dies in about five billion years from now, we will have further evolved into creatures very like us but sufficiently different to constitute a new species. We may hope that our successors, these new species, will be like us but hopefully better than us—better, more intelligent, more resilient, more able to cope with the new world that will exist then. They will either have to find us a new planet to live on, or they will have to develop a technology that will enable them to construct such a planet for themselves. It's a long time in the future, but we're going to face it, and I believe that the best insurance policy we have against these ultimate dangers, as Julian calls them, is the acceleration of science, not the retardation of science, and I believe the only way we will achieve that is if we can improve our cognitive powers, our powers of thought. So I think our salvation lies in beefing up cognitive enhancement and not worrying too much about moral enhancement, because there's the cognitive enhancement that will help us produce the science that will help us better confront the dangers that we face.

Questioner: So that was one sentence?

Harris: I specialize in long sentences.

Questioner: So we came to the conclusion that scientific progress is both the means of our salvation as well as the means of our downfall. Professor Savulescu, please tell us something about this hypothetical "God machine."

Savulescu: Some of these interventions to make people more moral, to make them better people, may affect their freedom. Some of you will have seen the film *A Clockwork Orange*, which was a very famous depiction of trying to stop violence through a form of essentially

brainwashing or conditioning. Surgery was performed in the 1950s and ’60s to stop people behaving immorally. Now, many people think of these with great horror, and I do too, and it’s possible that some kinds of enhancements won’t raise these sorts of issues. So, for example, if you make children more empathetic, for example, you don’t reduce their freedom; you just make them more like the people who are already empathetic who presumably are as free as anyone. Women are less violent and more empathetic. So if you made men more like women, as in part you do with pedophilia when you offer hormonal castration, you don’t reduce people’s freedom, because women are clearly as free as men. So these sorts of interventions won’t affect people’s freedom, but more radical interventions might. For example, if you had a greater understanding of psychopathy and were able to rewire the brain so that people thought in different ways or behaved in different ways, this would be something that removes people’s freedom, but so does imprisonment, and so do surveillance cameras, just as wearing a seatbelt. There are many things in society that restrict our freedom, to stop us from harming other people or, in some cases, ourselves. The God machine was an extreme example of this which I created to illustrate this. It is a science fiction example.

Imagine that we were able to construct in 2045 a supercomputer that was able to get information from all around the world and was also able to read people’s minds, and if they formed the desire to kill someone or commit a very violent crime, the God machine would immediately intervene and change their mind. However, the people just thought they changed their mind, but in fact it was this machine that intervened to stop murder, for example. This would remove people’s freedom to act very immorally. In my view, it might be worth it, because most of us would never form that intention to kill, so our lives would go on exactly as they do today. For those people who did form the intention to murder, if they did act on it, they would have a victim and would also end up in prison for the rest of their life, or even executed in the United States. To change their intention to kill is better than those sorts of things. In some cases, moral enhancement may restrict freedom, but I think it’s an open question whether that price is worth paying.

Questioner: Professor Harris, do you think that a God machine is the ultimate solution or it would be an act of selling ourselves into slavery?

Harris: I do think it would be an extreme folly to produce a God machine, for a number of reasons. The first is its catastrophic effect on freedom. This machine is going to be able—it’s a science fiction example that may never be possible—but we have to suppose that it would be, and then the question is, Would we want this machine to rule over our lives or not? What it does is to detect when you have a murderous thought, let us say, and it changes your mind for you so you no longer have that thought. You decide not to murder the person you were about to murder, and that’s allegedly a moral improvement. There are two problems with this. The first is that it clearly takes away our freedom. My own view is that freedom consists not only of the freedom to make mistakes but actually to do wrong.

Questioner: So without freedom to fall, good cannot be a choice.

Harris: Yes, “the freedom to fall” is a lovely phrase that I get from the late John Milton in his famous poem *Paradise Lost*, he imagines God talking to man and saying, “I made you

sufficient to have stood—that is, sufficient, able to make up your own mind but free to fall." Sufficient to stand but free to fall. And that is, as it were, a description of the alleged condition in which God made human beings, and the sufficiency to stand, the ability to make up our own mind, is the mirror image of the freedom to fall. It's not your own mind if you can't make it up and carry through that resolution, so I believe this machine would be the end of freedom. I also have another problem with this machine, and it goes back to what I was saying earlier about the very primitive nature of these dispositions that moral enhancement might operate on. You use or adapted my example of you rescuing, let us say, your daughter from a murderer, from somebody who was going to murder or rape her, by killing him. And so you intended to murder him. How would the God machine differentiate?

Savulescu: Oh, it would see the circumstances. It would see who the attacker is.

Questioner: So the machine will have a possibility and accessibility to context?

Savulescu: Yes. The idea is that it has perfect information about the world and your mental states—so it's not just reading your mental states. We see somebody that's about to murder this child, the machine would be able to interpret those actions, and it would change the murderer's mind. Think about this. We're all obsessed, rightly, with terrorism. If we have extremely effective antiterrorism programs that constantly thwarted terrorism, it seems little difference to me to use such a machine to pick up and change the intentions of terrorists. I don't think the freedom to commit terrorist attacks is a freedom worth protecting. Freedom without value is meaningless.

Harris: I'm delighted that Julian is so confident about his own inventions even where they're inventions of the imagination.

Savulescu: It's a philosophical discussion.

Harris: But I am deeply skeptical about having that degree of faith in machinery on which we would rely for the protection of our children, on which we would rely to allow us sufficient freedom to create our own lives. It's a thought, Julian.

Savulescu: We're dependent on machines all the time. We're actually going to create driverless cars. They will be safer.

Harris: You may be going to create, but I certainly am not.

Savulescu: That will take away your freedom to drive your car, but what would be the problem if the state said, "Well, these are so good we're not going to let people drive their cars; the car is going to drive itself." It would be not public transport but personalized transport.

Harris: Because it wouldn't be fun. I love driving.

Savulescu: It will happen.

Harris: I have no interest at all in a driverless car. I can't imagine anything more terrible. But to go back to the God machine, which is not my creation, so I don't have to live with the consequences—I wouldn't trust it an inch, and if it were as hyperintelligent as it would have to be to read everybody's thoughts in a millisecond quick enough to change their mind, what machine like that would ever give up that power? Part of Julian's example, which we haven't discussed, is he imagines us adults will voluntarily be able to sign up to this machine. So they wouldn't have to be plugged in to it if they didn't want to be; but once they're plugged in, it would effectively have denied them freedom and their future character—all of our character—which we build ourselves from the decisions that we made. We are all a product of our past decisions. They form our character. They make us the sort of people we are. I'm the sort of person who likes driving cars. All of these decisions to acquire tastes, habits, are what makes me the individual I am. Allow the machine to control even a part of your decisionmaking, and of course you no longer learn from your mistakes, because it won't let you make moral mistakes, and you've become the creature of that machine.

Savulescu: But this is nonsense. It's all-knowing and restricted to only affecting one kind of decision.

Harris: It'd be kind of you to let me complete the thought.

Savulescu: But you're going on for so long.

Questioner: Don't worry, you'll have your five minutes.

Savulescu: Your car does not allow you to skid now. You don't have the freedom to skid because it has an antilock brake system.

Harris: I can switch the traction control off in my car, and I do because it's more fun without it and I don't want—

Savulescu: You shouldn't trust it. Who knows? It might stop at the wrong time. You shouldn't trust those machines [laughing].

Harris: Your car wouldn't trust you, because you're a crap driver, but it would certainly trust me.

Savulescu: I never realized you were such a Luddite.

Harris: But to get back to the point, this God machine would effectively take over our minds and our personality, and anybody who would trust a machine that intelligent to ever let us go again, (a) it wouldn't, and (b) we wouldn't want to be the people who would want to be let go. We wouldn't notice anything. It would change our mind without us knowing our mind is being changed. So if you plug yourself into the machine for a year, and then you unplug yourself and you're asked, "Do you want to go back in again?" you would say, "Well why not? It hasn't made any difference. I've done everything I wanted to do, so what's the problem?" But of course, you don't know what you would've wanted to do. It would've stopped you doing. It's a nice example, but it seems to me too much of a surrender to a totalitarian nightmare of the brave-new-world sort, and, if it is as voluntary as Julian implies,

I am skeptical about its utility; namely, that only competent adults—children might be compulsorily plugged in—would have to choose to be plugged in. I don't think any sane, competent adults would allow themselves to be plugged in, but of course if they had, they'll never have a motive to be unplugged again. It's basically selling yourself into slavery. It may be a benign slavery, but I doubt it.

Savulescu: I think it's important to recognize that many very respectable scientists and philosophers believe that free will doesn't exist at all. Now, I'm inclined to think that we do have free will, but I think it's much smaller than you think it is. John is going on about totalitarian machines controlling your actions. Your actions are determined by influences outside of your control every single day. If you put a human eye above a tin for coffee money, it makes people more honest. Even three dots in the shape of a face will make you more honest than you would be if they weren't there, or if they were inverted. If you're hungry, you're more severe in your judgments. As a judge, if your desk is dirty, you are more likely to be severe and imprison people for longer sentences. So there are so many factors that influence our judgments, including the actions of other people, society, and so on. This idea that we're a great free agent choosing whether we act or not is a fiction. There's a strong argument that many of the psychopaths who commit these sorts of violent crimes are genetically disposed. It's not as if they just choose freely to kill someone; it's largely out of their hands. That is the theme of the Coen brothers film, *No Country for Old Men.* How much out of their hands? We don't know. But I think this appeal to freedom is highly overrated.

Harris: If the God machine does remove freedom, you will never know that, because it would have changed your mind without you noticing it. You'll never know the extent to which it may have intervened in your life. It's like censorship. I am against censorship. If you allow a censor, somebody who stops the printing or the publication or the putting on television or on radio of ideas that it judges too dangerous for citizens to be exposed to, then you lose control – you never know when you have been the recipient of censored material.

The censor can be a machine, it could be a person, you can give them a remit, you can give them instructions to only censor the really naughty things, the things that would be corrupting for us to see, but once you have censorship, you never know what you're not being allowed to see. And so you never know whether the censor is functioning in the way that it is supposed to function. Once you sell yourself down that river of censorship or allowing somebody to intervene between you and your freedom, whether it's freedom to read or to be exposed to ideas or whether it's freedom to decide, you never know anymore. You've abdicated responsibility for your life once and for all.

Savulescu: This is a great example. I think it's a real open question whether the free access to hardcore pornography on the Internet is better or worse for people, and indeed whether it's ruining our children's lives. I don't think there's a straightforward freedom right to access hardcore pornography on the Internet. I think it should be an open question whether we impose constraints on access to pornography. But let me ask you: let's say the God machine does work perfectly. It's not going to change any other intentions. It's not going to be used for other purposes. It's just there. We have, as a society, a choice whether to

implement its use to affect everyone in our society. Hands up, those of you who would prefer to prevent terrorism and murder by the use of the God machine. Anyone? There are some rational people out there. The rest may need cognitive and moral enhancement.

Questioner: Who are the guardians of the guardians? And who should decide what qualities people ought to have. Anyone?

Savulescu: So how do you decide which qualities to aim for? Again, I make this point, and this gets back to your point about who are the guardians. We today, we have guardians, and we are the guardians of the guardians. We decide whether there's a death penalty. We decide how long to imprison people for. We decide what sort of antiterrorism measures to employ. We're making these decisions all the time. Likewise, we're making decisions about what kind of moral values we want our children to have, our society to have. Now, the question is not, Can we avoid this? The question is, What should we really be aiming for? and it's my claim that we should aim for an ethic that promotes the values of justice, tolerance, respect for human equality, a sense of altruism, and willingness to cooperate and make small self-sacrifices for the benefit of others. Now, you can disagree, and we can have a discussion, but you can't escape commitment to some kinds of values. This is the sort of thing that will be instilled in society anyway. And it's not clear to me that we have the best set of values now, and it's not clear to me that we should only aim to enhance through social means.

Harris: And it never will be clear to you or to me, whether we have the best set of values. That's why they have to be constantly revisable, and that's why we need to be aware of the decisions that are being made, and we need to learn from our mistakes. If we can't learn from our mistakes, we don't even know that there are mistakes. The God machine will prevent you even thinking what you've done was a mistake.

Savulescu: This was a hypothetical example. If you think that learning from your mistakes is important, then people will differ in that respect, and they're going to enhance that.

Harris: If you're telling us you chose a very bad example, I would agree with you.

Savulescu: This example was intended to show that freedom has a limited—not unlimited—value, as John thinks it has. Freedom to do certain things is not the sort of thing that we need. What we actually decide to enhance is a question of great significance and a question that we have to answer. Some people say to me, "Look, you should just be tolerant of everyone." But people differ in the degrees to which they're tolerant, and you can change that. You can change the ability to learn from their crimes or mistakes or so on. The question is—and this is so for all technology, with the power to change things—how should it be used? With technological power comes responsibility to make a decision. To choose to do nothing, to leave the status quo, is something you're responsible for and the consequences of that. Now, you can't absolve yourself of responsibility later, and we face difficult decisions whatever we decide.

Questioner: Would many of us be really motivated to reach out for moral enhancement? In modern society, people with high moral standards, strong empathy, and willingness to cooperate are often abused.

Savulescu: This is a good example that moral enhancement needs to be a wide concept with many different dimensions. It's not the Christian doctrine of "turn the other cheek." If you turn the other cheek constantly, you will be exploited. So that isn't a moral enhancement, as John said at the beginning. If it's not good for you, it's not an enhancement. If it doesn't promote a moral outcome, it's not a moral enhancement. So, if it's the case that being too altruistic will lead to exploitation and harm to you, then don't increase altruism. But there will be, as John says, cognitive enhancement, which is a very important aspect of moral enhancement. What other things should be in the mix? I think we can study that and look at what sorts of things will deal with the problems we face today. I don't know what the correct level of empathy is. I don't know what the correct level of cooperativeness is. What I do know is we should be trying to answer those questions systematically and not just sticking our head in the sand like ostriches and saying, "Human beings, they're good enough," or, "We shouldn't use medicine or science to change human beings."

Harris: I know that Julian would agree with me that I've never suggested anything like that. In fact, our views are very similar. We both believe strongly that our world and ourselves face huge dangers, possibly insoluble dangers. We both agree, I think, that we need to be open to new and radical solutions to those dangers and safeguard against them. We disagree about a small corner of the methodologies for doing that. I think it's important to emphasize that the range of our agreement, both about the problems facing us and about the types of solutions that we need to find to those problems, is enormous. The disagreement is enormous about only some of the chosen methods. But that's small compared with our agreement. I think you've seen that we're quite capable of disagreeing quite strongly with one another but I think the things on which we agree are much more important and much more fundamental. We agree on the nature of problems we face. We agree about their seriousness. We agree that radical new and courageous solutions are required to solve those problems. We have some disagreements about the respective merits of some of those possible solutions, but we're also agreeing about something very important, and that is that the search for those solutions has to go on. I'm not against looking for methods of moral enhancement, but I just don't think we've found any promising ones yet, but when, and if I do think we've found promising ones, I will support them in the way I support 90 percent of what Julian also supports in terms of human enhancement, in terms of making our world a better and safer place for all of us and preparing for a future which is going to be much, much more difficult than the past.

Questioner: A question from the audience for Professor Harris, "I'm a student of philosophy here. You said that if we don't have right to fail, if you don't have a right to make a moral mistake, you wouldn't learn, but why do we have to learn if we can't make a mistake? I mean, the point of learning is to stop making mistakes."

Harris: Yes, well, what constitutes a mistake is a moveable feast. Mistakes are not necessarily mistakes once and for all; many things that were regarded as mistakes in the past are no longer regarded as mistakes. So we don't want to mortgage our future to a

particular approach to our problems. We need to keep an open mind and open abilities. I think it's naïve to think that we can eradicate mistakes. It would be nice, but I think our conception of what will constitute a mistake will change, as it has changed over time. Many people were very sanguine, very complacent about the introduction of the motorcar. And much as I love driving, it's a very mixed blessing. I can also confidently predict there will be no internal-combustion-engine-powered cars in the future, much as I am a petrol-head and love them. So yes, if we could be sure that we would never in the future ever make any mistakes, and be very clear about what constitutes a mistake, that would be wonderful. I just have no confidence that we can make those discriminations either now or in the future with sufficient reliability to sign up to something like the God machine or it's perhaps more realistic equivalent.

Savulescu: The question was a very good one. The stakes have changed, because some mistakes today will be so catastrophic, there will be no learning from them. Take the example that John gave of publishing the way to engineer an infectious version of bird flu. Now, it's always served us well to have scientific freedom to publish everything and make it widely available. It's not clear to me that when the harm arising from such freedom is bioterrorism, an intentionally created pandemic killing twenty million people, that we shouldn't restrict that information to people who have a security clearance (which already does happen in the employment of people in highly infectious or dangerous laboratories). So the mistakes that we can make today, because of the power that we have, are completely different from the mistakes that we've made through human history that we can afford to learn from. That's the point of this discussion: that we are in a different phase of human existence than we've ever been in before. We can't simply say, "Well, scientific freedom—that's served us very well in the past, so it will in the future." It won't necessarily do for the future.

Audience member: So it seems to me that moral enhancement, as you present it, must be universal in order to function. If it is not universal and not compulsory, how can it work, because I guess we can assume that those people who choose to be morally enhanced are already those who are reluctant to act immorally, and if it is universal, you might say it's heaven on earth. On the other hand, it just could be seen as a way to use our freedom to destroy our freedom, which might be okay for us, like if you committed suicide, but is it something we can do for the future generations?

Savulescu: It's never going to be universal. You're never going to do something that affects seven billion people. All you'll ever be able to do is change probabilities. You can't obtain certainty anywhere in the world, so how might we use this in a way that has some effect, some beneficial effect? Our key political leaders make hugely important decisions. Imagine if Al Gore had got in as president instead of George W. Bush. We may not have had a war on terror. We may have a completely different world. It may be that moral enhancement changing or influencing the moral environment of the key decisions that are made or peace negotiations or whatever, is a way of significantly shifting the probabilities. If we could make progress on the conflict between Israel and Palestine, that would be a wonderful thing. These are moral decisions that are made. Now, we already changed the environment to maximize certain outcomes in negotiation or political decisionmaking. Why should we leave the major decisions up to extrinsic factors like how dirty the desk is or how tired the person

is? And these are things that we developed a science of. Maybe we can't have any effect on anything, but the point is that we should try.

Perhaps you can enhance the acquisition of moral skills in children in developed societies. That's not going to cure the world's problems, but it might make some progress. We will never get rid of these risks now that we've created powerful technology for ourselves. But the best we can do is to try to reduce them. People have spent a huge amount of money and time talking about social reform; these are important and maybe they'll provide the answers. However, modest progress has been made—perhaps it is time that we should also look inside ourselves for the causes of these problems. Not because we believe they're going to generate solutions anytime soon; but so significant are the problems we face today, we should explore every avenue. I personally am very pessimistic. I'll leave you with this thought. Stephen Hawking, the famous physicist, did a poll a number of years ago, asking people what they thought the biggest risks were this century. At the end, they asked Hawking what he judged were the biggest risks for the survival of the species. He said: nuclear war, climate change, natural epidemics, bioterrorism, these sorts of things. He said it will take at least one hundred years for humans to be able to start to inhabit space, as John has suggested, that we will need to ensure the long-term survival of the species, and that this is a critical century. Martin Rees, the president of the Royal Society in England, called this our "final century." Hawking finished his interview by saying that, in order to survive, perhaps we need to genetically enhance ourselves to make ourselves wiser and less aggressive. Now, I don't know whether it's going to come by genetic intervention, but it's certainly an issue that people like Stephen Hawking and Martin Rees see as a very important project for us to discuss and pursue. I don't know if we can make ourselves wiser and less aggressive, but I think we should at least research it.

Harris: I agree.