Moral enhancement

Routledge Encyclopaedia of Philosophy

Lisa Forsberg and Thomas Douglas

Article summary

Moral enhancements aim to morally improve a person, for example by increasing the frequency with which an individual does the right thing or acts from the right motives. Most of the applied ethics literature on moral enhancement focuses on moral *bio*enhancement—moral enhancement pursued through biomedical means—and considers examples such as the use of drugs to diminish aggression, suppress implicit racial biases, or amplify empathy.

A number of authors have defended the voluntary pursuit of moral bioenhancement, or the development of technologies that would enable it. They have highlighted the need for humans to morally improve themselves in order to address moral failures such as the oppression of women, the mistreatment of animals, and anthropogenic climate change. They have also emphasised the moral similarities between moral bioenhancement and more familiar forms of moral enhancement, such as that achieved through childhood education, introspective reflection, and engagement with literature.

Critics of moral enhancement have argued that it may undermine our freedom to 'fall' (i.e. be immoral) and therefore our moral agency, or exacerbate the domination of individuals by political authorities. They have also questioned the potential for biomedical interventions to produce the deepest and most valuable forms of moral improvement, and have highlighted the risks that technologies for moral bioenhancement might misfire or be intentionally misused, thereby producing moral deterioration. Underlying some of these worries is the observation that there is little agreement on which psychological transformations would constitute moral improvements, and in which contexts. Defenders of moral enhancement have made various proposals for resolving or side-stepping these disagreements, but it remains unclear how far these proposals can take us beyond establishing consensus on the worst types of moral failure.

1. What is moral enhancement?

Moral enhancements are interventions that aim to, and can be expected to succeed in, morally improving a person. They can be categorised according to their target, their metric, and their means. The target is the psychological or behavioural feature that the user of a moral enhancement seeks to improve. Targets of moral enhancement could include, for example, patterns of conduct (for example, a tendency to act violently), emotional capacities (for example, the capacity for empathy), or character traits (for example, aggressiveness). The metric of a moral enhancement is the dimension on which moral improvement is sought. Dimensions might include, for example, moral permissibility, desirability, or praiseworthiness. Thus, for example, a moral enhancement might seek to increase the frequency with which an individual performs morally desirable acts, such as charitable giving to worthwhile causes, or decrease the frequency of morally impermissible acts, such as acts of gratuitous violence). The means of a moral enhancement are the processes via which the user seeks to increase the value of the target. Philosophers often distinguish biomedical means, such as the administration of drugs or the

modification of a person's genetic make-up, and environmental means, such as education programmes and counselling.

Most of the applied ethics literature on moral enhancement focuses on moral *bio*enhancement—moral enhancement pursued through biomedical means—and considers examples such as the use of drugs to diminish aggression, suppress implicit racial biases, or amplify empathy. It examines when, if ever, it would be morally permissible, desirable or obligatory to (a) seek to develop such interventions, (b) voluntarily undergo them oneself, and (c) impose them on others.

2. Arguments for moral bioenhancement

A number of authors have defended voluntary biomedical enhancement, or the development of technologies that would enable it.

Thomas Douglas (2008) argues that it would often be morally permissible for an agent to voluntarily alter herself in ways that could reasonably be expected to morally improve her future motives. He notes that many of the objections to other forms of biomedical enhancement—such as those targeting physical or cognitive abilities—focus on potential negative effects such interventions might have on individuals other than the recipient (for example by leading to unfairness or exacerbating inequality). These objections, Douglas argues, do not apply to interventions aimed at making recipients more moral, which would typically *benefit* unenhanced individuals (cf Archer 2016).

David DeGrazia (2014) argues that traditional means to moral improvement, such as moral education and mentoring, socialisation, and literature encouraging reflection, are insufficient to bring about the degree of moral improvement needed to avoid serious moral failures—for example, to alleviate and prevent poverty, oppression and violence against women, and mistreatment of non-human animals. This gives us reason to seriously and open-mindedly consider moral bioenhancement. But in doing so, we face a serious challenge: people hold different views about what morality requires and about how we ought to act. Given moral disagreement, how can we determine what changes should be accepted as moral enhancements? DeGrazia's solution is to appeal to the same strategies that we employ when pursuing traditional, environmental forms of moral enhancement, such as the moral education of our children: to focus on changes that are acceptable from any reasonable moral view. Such changes might include improved ability to acknowledge unpleasant realities (e.g., 'starving children, the abuse of women, the worst conditions of factory farms') and attenuation of violent impulses. Provided that moral bioenhancements were safe, effective, and could be made available to those who might not otherwise be able to afford them, their use for these ends would, DeGrazia argues, be desirable.

Ingmar Persson and Julian Savulescu (2012) start from the premise that human moral psychology developed in, and is therefore best adapted to, societies that were smaller and more close-knit than the societies we live in today. Technologies to which we had access were primitive and inefficient and allowed us only to affect our most immediate environment. Science and technology have radically transformed our living conditions, but human moral psychology has not kept up. We therefore now find ourselves in a situation in which we have at our disposal technologies that allow us to exert influence 'all over the world' and 'far into the future', but without having the moral capacities to handle the moral problems such technologies and our

way of life moral generally raises—problems such as climate change and the threat of bioterrorism. There is therefore, Persson and Savulescu argue, 'an urgent imperative to enhance the moral character of humanity'. Since traditional means by which we might pursue moral improvement are unlikely to bring about the changes needed, we ought to pursue the development of moral bioenhancements which could serve as a complement to these means.

3. Objections to successful moral bioenhancement

Objections to moral bioenhancement may be divided into three broad categories. Objections in the first category question the conceptual possibility or technical feasibility of bringing about moral improvement through biomedical means (for discussion, see, for example, De Melo-Martín 2018). Objections in the second category allow that biomedical technologies could produce moral improvement, but deny that they are likely to (reliably) do so (see §4). Objections in the third category hold that, even if moral bioenhancement would be successful—that is, would produce moral improvement—it would be morally impermissible or undesirable to pursue it (this section).

3.1 Restriction of freedom

Some worry that moral bioenhancement—or certain types of moral bioenhancement—may restrict the enhanced individual's freedom, autonomy, or moral agency in ways that are problematic. For example, Harris (2011, see also Ehni and Aurenque, 2012; Huang 2018; Sorensen 2014) argues that non-cognitive moral bioenhancements—those which operate via the direct (psychologically unmediated) modulation of emotional or motivational states—threaten our moral agency, by denying us the option of acting immorally or by denying us opportunities for self-governance. Hauskeller (2017, see also Harris 2011; 2016a) argues that a world in which individuals have moral freedom and occasionally perform morally undesirable acts is better than a world in which individuals perform (only) morally desirable acts but lack this freedom, because the ability to choose between moral and immoral is the 'essence of being human' and of 'far greater intrinsic value' than performing morally desirable acts.

This 'freedom objection' to moral enhancement has been met with two main types of response. The first holds that even when moral enhancement does restrict freedom, the value of that freedom will sometimes be outweighed by other values (see e.g. DeGrazia, 2014; Douglas, 2008; Persson and Savulescu, 2016). Restricting one's future freedom, for example by freely entering into a 'Ulysses contract' that restricts one's future conduct, is often a morally desirable thing to do. The second response disputes the claim that *all* non-cognitive moral bioenhancements would restrict freedom, thus limiting the objection to a smaller subset of interventions. Some versions of the second response note that agents who have voluntarily pursued non-cognitive moral bioenhancements would seem to be no different, in their freedom, to individuals who are 'naturally' particularly virtuous, and we would not generally think that the latter individuals have less freedom than others. Other versions of this response hold that some non-cognitive moral bioenhancements might operate by increasing an individual's freedom to be moral, rather than reducing her freedom to be immoral (see e.g. Douglas 2008; Earp, Douglas, Savulescu 2017). They might, for example, attenuate emotional impulses or influences that would otherwise constrain freedom.

Sparrow (2014a, 2014b) and Hauskeller (2017) both defend versions of the 'freedom objection' that focus not on *metaphysical freedom* but instead on *political freedom*. On these versions of the objection, the problem with moral bioenhancement concerns not the moral agency or responsibility of the enhanced individuals (as in the metaphysical version), but asymmetries of power—'how others relate to and express their wills over us' (Hauskeller 2017, see also Habermas 2003, and Sparrow 2014a, 2014b). On Sparrow's version, the worry is about state-sanctioned moral bioenhancement, which he claims would shape us in ways that do not allow us to conceive of our own moral ideals and identity (Sparrow 2014a, 2014b, see also Shaefer 2015). But Hauskeller (2017) suggests that even voluntary self-enhancement could raise this worry, threatening the ability of the enhanced individuals to see themselves as the moral equals of others. Danaher (2018) responds to the political freedom objection by arguing that our moral behaviour is influenced by many factors, including genetic, social, environmental, and educational factors, and that it is not clear why moral bioenhancement should be singled out as particularly problematic from the point of view of freedom. Similarly, Huang maintains that Sparrow's version of the objection would exclude many forms of moral education we accept and endorse (Huang 2020).

3.2 Superficiality

Another objection to moral bioenhancement is grounded in what we might call the 'superficiality concern': the idea that while such enhancement might successfully induce morally desirable conduct, it would not produce the deeper kind of moral improvement that we should be seeking. The morally desirable behaviour resulting from moral bioenhancement would somehow lack full moral value or worth. One underlying idea seems to be that because traditional means operate via the agent's understanding of what morality requires and why, these means will produce more stable and reliable morally desirable dispositions, while the moral improvement brought about via moral bioenhancement might be 'accidental' and therefore less likely to be stable and reliable (for concerns of this kind, see Harris and Chan 2010; Harris 2012a, 2012b, 2016a; Jotterand 2011; Sparrow 2014c).

It might be responded that even if this is true, it does not follow that moral bioenhancement is morally undesirable or impermissible, only that it is in one way inferior to traditional forms of moral enhancement. Some have also argued that moral bioenhancemeents might operate precisely by promoting moral understanding, for example by attenuating emotions that serve as barriers to it (Douglas 2008) or promoting 'learning through doing' (Wasserman 2011).

4. Misuse or misfiring

Persson and Savulescu's defence of moral bioenhancement is partly motivated by the concern that powerful technologies might be used in harmful ways. But, opponents note, the same concern could be directed at technologies used for moral bioenhancement (see e.g. Ehni & Aurenque 2012; Harris 2011). The biotechnologies that would enable moral bioenhancement could presumably also be used to produce moral *dis*enhancement. Opponents of moral bioenhancement also point out that the harmful events of concern to Persson and Savulescu, such as major bioterrorist attacks, could be caused by a lone malevolent agent, and such agents are unlikely to choose to undergo moral bioenhancement, so the availability of such enhancements will do little to reduce the risk of harmful events (Harris, 2011). In response, Persson and Savulescu (2012) argue that technological progress increases the risk not only of

lone agents causing harm, but also of great harm occurring as a result of more mundane moral failings dispersed across a larger number of people (e.g. contributions to climate change). If moral enhancement can reduce the rate of moral failure among some proportion of the population, then, it can at least reduce *some* risks of great harm, even if not used universally.

Others worry that even well-intended attempted moral bioenhancements might misfire, causing moral deterioration rather than improvement. For example, empathy enhancement might make recipients too partial to their loved ones or their in-group. The risk of misfiring might be especially high for individuals whose societal functions or professional roles require traits that would ordinarily be morally undesirable (Wasserman 2011). For example, experiencing more empathy might interfere with a person's performance as a neurosurgeon or judge even if it would be morally desirable in her private life.

Harris (2011, see also Jotterand, 2011; Wasserman, 2011) argues that misfiring is likely since technologies for moral bioenhancement will be blunt, while the changes required to bring about morally desirable behaviour are complex and fine-grained (see Agar 2010; 2015). In response, proponents of moral enhancement concede that complex and fine-grained interventions will often be needed, but dispute that moral enhancements will be incapable of achieving fine-grained alterations (Douglas, 2013b; Earp, Douglas, Savulescu 2017). Douglas (2011) notes that, in clinical psychiatry, blunt instruments are often used when complex and fine-grained changes are needed. There, the risk of negative effects can often be effectively mitigated by close monitoring and action taken if adverse effects are detected. A similar approach could be taken in moral bioenhancement.

Some worries about misfiring derive from empirical uncertainty about what the effects of a given biomedical intervention will be. But others derive from moral uncertainty—uncertainty about which kinds of psychological or behavioural changes qualify as moral improvements. For example, concerns have been expressed about whether reasonable disagreements about the proper targets of moral bioenhancement—for example, about which motives and conduct moral bioenhancements should seek to prevent or foster—could be satisfactorily resolved. If they could not, we might worry that many will pursue what they reasonably take to be, but are in fact not, moral improvements. One response to the concern about disagreement has been to avoid committing to any particular view of moral goodness, and what determines it (e.g. Douglas 2008). But this leaves us without any practical guidance on which moral bioenhancements, if any, to pursue. A second response to the concern about disagreement is to focus on moral improvements on which there would be a broad consensus. DeGrazia (2014) argues that while there are a wide range of moral views that contradict each other on many matters, there are nevertheless 'points of overlapping consensus among competing, reasonable moral perspectives'. We can turn to traditional moral education (for example of children or criminal offenders) to identify changes that plausibly lie within this consensus; DeGrazia suggests, as promising candidates, the attenuation of 'sadism, … defective empathy, out-group prejudice, inability to face unpleasant realities, weak will, impulsivity, lack of nuance in moral understanding, and inability to compromise'. Alternatively, we might prefer moral bioenhancements that target psychological *processes*, such as rational deliberation, that generally promote agents' moral reliability. We might do this by, for example, enhancing 'logical competence, conceptual understanding, empirical competence, openness, empathy' and unbiasedness (Schaefer and Savulescu 2019). Such *procedural* moral bioenhancements may preserve what we value in moral disagreement. For example, it might allow us to morally improve ourselves without sacrificing our distinctive

individual moral commitments or locking ourselves into moral views about which we are uncertain (Schaefer 2015; Schaefer and Savulescu 2019).

5. Further questions

Existing discussion of moral bioenhancement raises at least as many questions as it answers. Some of these are fundamental questions in moral theory: What does it take to be moral or act morally? What is the relative importance of doing the right thing, and doing it for the right reasons? And does the value of moral improvement depend more on the journey or on the destination? For example, is moral improvement more valuable when it is difficult to achieve, or should we prefer less effortful means to improvement, other things being equal? The possibility of moral bioenhancement forces us to confront these questions because it raises the prospect of having to trade-off different types of moral improvement against one another. For example, we may face choices between biomedical means to moral improvement, which require less effort and are more effective at producing right action, and traditional means, which are better at producing deep moral understanding and sensitivity.

Questions also remain regarding the extent to which moral disagreement and moral uncertainty pose a problem for moral (bio)enhancement. Though proponents of moral bioenhancement might be right to think that consensus could be reached on some less ambitious forms of moral enhancement (e.g., reducing impulsive violence), consensus will be harder to achieve when we are thinking about radical changes that might be required to avoid existential threats (e.g. due to climate change), if we aim to achieve morally permissible treatment of non-human animals globally, or when we are thinking beyond currently existing people. Such changes might require more radical and less widely agreed upon changes, and difficult choices between beneficiaries (e.g. current or future generations).

Discussion has focussed on whether technologies for moral bioenhancement ought to be developed, and whether individuals ought to voluntarily use them. But there are also important questions about, for example, whether and when healthcare professionals ought to *offer* moral bioenhancements, and whether the state ought ever to *require* their use, for example, in the context of criminal rehabilitation, perhaps as a condition of release from prison. The latter raises questions concerning whether consent to undergo such moral bioenhancement is valid when the alternative to consenting is further imprisonment, and whether interventions that interfere with recipients' bodies and minds may be permissible when administered without recipient consent (Bomann-Larsen 2011; Rosati 1994; Douglas 2014c).

6. Bibliography and further reading

Agar, N. (2010) 'Enhancing genetic virtue?' *Politics and the Life Sciences* 29(1): 73–75.

Agar, N. (2015) 'Moral Bioenhancement and the Utilitarian Catastrophe, *Cambridge Quarterly of Healthcare Ethics* 24: 37–47.

Archer, A. (2016) 'Moral Enhancement and Those Left Behind', *Bioethics* 30: 500-510.

Baertschi, B. (2014) 'Neuromodulation in the Service of Moral Enhancement', *Brain Topography* 27: 63–71.

Bomann-Larsen, L. (2013) 'Voluntary Rehabilitation? On Neurotechnological Behavioural Treatment, Valid Consent and (In)appropriate Offers', *Neuroethics* 6: 65–77.

Bublitz, J.C., Merkel, R. (2014) 'Crimes Against Minds: On Mental Manipulations, Harms and a Human Right to Mental Self-Determination, *Criminal Law and Philosophy* 8(1): 51–77.

Bublitz, J.C. (2016) 'Moral Enhancement and Mental Freedom', *Journal of Applied Philosophy* 33(1): 88–106.

Bublitz, J.C. (2018) '"The Soul is the Prison of the Body": Mandatory Moral Enhancement, Punishment and Rights Against Neuro-Rehabilitation', in Birks, D. and Douglas, T. (eds). *Treatment for Crime: Philosophical Essays on Neurointerventions in Criminal Justice*, Oxford: Oxford University Press: 289-320.

Carter S. (2017) 'Could Moral Enhancement Interventions be Medically Indicated?' *Health Care Analysis* 25(4): 338-353.

Danaher, J. (2018) 'Moral Enhancement and Moral Freedom: A Critique of the Little Alex Problem', *Royal Institute of Philosophy Supplement* 83: 233-250.

DeGrazia D. (2014) 'Moral Enhancement, Freedom, and What We (Should) Value in Moral Behaviour, *Journal of Medical Ethics* 40: 361-368.

De Melo-Martín, I. (2018) 'The Trouble With Moral Enhancement', *Royal Institute of Philosophy Supplement*, 83: 19-33.

Douglas, T. (2008) 'Moral Enhancement', *Journal of Applied Philosophy* 25(3): 228–245.

Douglas T. (2013a) 'Moral Enhancement via Direct Emotion Modulation: A Reply to John Harris', *Bioethics* 27(3): 160–168.

Douglas, T. (2014a) 'Enhancing Moral Conformity and Enhancing Moral Worth', *Neuroethics* 7: 75–91.

Douglas, T. (2014b) 'The Relationship Between Effort and Moral Worth: Three Amendments to Sorensen's Model', *Ethical Theory and Moral Practice* 17: 325–334.

Douglas, T. (2014) Criminal Rehabilitation Through Medical Intervention: Moral Liability and the Right to Bodily Integrity, *Journal of Ethics* 18(2): 101–122.

Douglas, T. (2015) 'The Morality of Moral Neuroenhancement', in Clausen, J. and Levy, N. (eds.). *Handbook of Neuroethics*, Dordrecht: Springer: 1227–1249.

Douglas, T. (2018) 'Neural and Environmental Modulation of Motivation. What's the Moral Difference?', in Birks, D. and Douglas, T. (eds.) *Treatment for Crime: Philosophical Essays on Neurointerventions in Criminal Justice*, Oxford: Oxford University Press: 208-223.

Earp, B.D., Douglas, T., Savulescu, J. (2017) 'Moral Neuroenhancement', in Syd, L., Johnson, M., Rommelfanger, K.S. (eds.), *The Routledge Handbook of Neuroethics*, Routledge.

Ehni, H-J., & Aurenque, D. (2012) 'On Moral Enhancement from a Habermasian Perspective', *Cambridge Quarterly of Healthcare Ethics* 21(02): 223–234.

Faust, H. S. (2008) 'Should We Select for Genetic Moral Enhancement? A Thought Experiment Using the MoralKinder (MK+) Haplotype', *Theoretical Medicine and Bioethics* 29(6): 397–416.

Forsberg, L., & Skelton, A. (2020) Achievement and Enhancement, *Canadian Journal of Philosophy* 50(3): 322-338.

Glannon, W. (2018) 'Moral Enhancement as a Collective Action Problem', *Royal Institute of Philosophy Supplements* 83: 59-85.

Glover, J. (1984) *What Sort of People Should There Be? Genetic Engineering, Brain Control and Their Impact on Our Future World*, Penguin Books.

Greely, H.T. (2008) 'Neuroscience and Criminal Justice: Not Responsibility but Treatment', *Kansas Law Review* 56: 1103-38.

Habermas, J. (2003) *The Future of Human Nature*, Cambridge: Polity Press.

Harris, J., & Chan, S. (2010) 'Moral Behavior is Not What it Seems, *Proceedings of the National Academy of Sciences* 107.50: E183.

Harris, J. (2011) 'Moral Enhancement and Freedom' *Bioethics* 25(2): 102–111.

Harris J. (2012) 'What It's Like to Be Good', *Cambridge Quarterly of Healthcare Ethics* 21(3): 293-305.

Harris, J. (2013) 'Ethics is for Bad Guys!' Putting the 'Moral' into Moral Enhancement, *Bioethics* 27: 169-173.

Harris, J. (2013) 'Moral progress and moral enhancement', *Bioethics* 27(5): 285–290.

Harris, J. (2016a) *How to Be Good: The Possibility of Moral Enhancement*, Oxford: Oxford University Press.

Harris, J. (2016b) 'Moral Blindness—The Gift of the God Machine'. *Neuroethics* 9:269–273.

Hauskeller, M. (2017) 'Is it Desirable to Be Able to Do the Undesirable? Moral Bioenhancement and the Little Alex Problem', *Cambridge Quarterly of Healthcare Ethics* 26(3): 365–376.

Huang, P-H. (2018) 'Moral Enhancement, Self-Governance, and Resistance', *Journal of Medicine and Philosophy* 43: 547–567.

Huang, P-H., Tsu, P. (2018) 'Biomedical Moral Enhancement in the Face of Moral Particularism', *Royal Institute of Philosophy Supplement*, 83: 189-208.

Huang, P-H. (2020) 'Who's afraid of perfectionist moral enhancement? A reply to Sparrow', *Bioethics*. 34: 865– 871.

Jotterand, F. (2011) '"Virtue Engineering" and Moral Agency: Will Post-Humans Still Need the Virtues?' *AJOB Neuroscience* 2(4): 3–9.

Jotterand, F., Levin, S.B. (2019) 'Moral Deficits, Moral Motivation and the Feasibility of Moral Bioenhancement', *Topoi* 38: 63–71.

Levy, N. (2019) 'Nudge, Nudge, Wink, Wink: Nudging is Giving Reasons', *Ergo* 6: 281-302.

Levy, N. (2017) 'Nudges in a post-truth world', *Journal of Medical Ethics* 43: 495–500.

Levy, N. (2018) 'Nudges to reason: not guilty', *Journal of Medical Ethics* 44: 723.

Levy, N., Douglas, T., Kahane, G., Terbeck, S., Cowen, P. J., Hewstone, M., & Savulescu, J. (2014) 'Are You Morally Modified? The Moral Effects of Widely Used Pharmaceuticals', *Philosophy, Psychiatry, & Psychology* 21(2): 111-125.

Matravers, M. (2018) 'The Importance of Context in Thinking About Crime-Preventing Neurointerventions', in Birks, D. and Douglas, T. (eds.) *Treatment for Crime: Philosophical Essays on Neurointerventions in Criminal Justice*, Oxford: Oxford University Press: 71-93.

McMahan, J. (2009) 'Cognitive Disability and Cognitive Enhancement' *Metaphilosophy* 40(3-4): 582–605.

Pacholczyk, A. (2011) 'Moral enhancement: What is it and do we want it?' *Law, Innovation and Technology* 3(2): 251–277.

Parens, E. (2013) 'The need for moral enhancement', *Philosophers' Magazine* 62(1): 114–117.

Paulo, N., Bublitz, J.C. (2019) 'How (not) to Argue For Moral Enhancement: Reflections on a Decade of Debate', *Topoi* 38: 95–109.

Persson, I., & 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): 162–177.

Persson, I., & Savulescu, J. (2010) 'Moral Transhumanism', *Journal of Medicine and Philosophy* 35(6): 656 –669.

Persson, I., Savulescu, J. (2012) *Unfit for the Future: The Need for Moral Enhancement*, Oxford: Oxford University Press.

Persson, I., Savulescu, J. (2013) 'Getting Moral Enhancement Right: The Desirability of Moral Bioenhancement', *Bioethics* 27: 124-131.

Persson, I., Savulescu, J. (2015) 'The Art of Misunderstanding Moral Bioenhancement, *Cambridge Quarterly of Healthcare Ethics* 24(1): 48–57.

Persson, I., Savulescu, J. (2016) 'Moral Bioenhancement, Freedom and Reason', *Neuroethics* 9(3): 263–268.

Rakic, V. (2014) Voluntary moral bioenhancement is a solution to Sparrow's concerns. *American Journal of Bioethics* 14(4): 37–38.

Rosati, C. S. (1994) 'A Study of Internal Punishment', *Wisconsin Law Review* 123: 123-170.

Rumbold, B. (2017) Tying oneself to the mast: One necessary cost to morally enhancing oneself biomedically, *Bioethics* 31(7): 543-551.

Ryberg, J., Petersen, T. S. (2013) 'Neurotechnological Behavioural Treatment of Criminal Offenders—A Comment on Bomann-Larsen', *Neuroethics* 6: 79-83.

Savulescu, J. (2009) 'The Human Prejudice and the Moral Status of Enhanced Beings: What Do We Owe the Gods', in Savulescu, J., Bostrom, N., *Human Enhancement*, Oxford: Oxford University Press.

Savulescu, J., Persson, I. (2012) 'Moral Enhancement, Freedom, and the God Machine', *The Monist* 95(3): 399–421.

Savulescu, J., Douglas, T., Persson, I. (2014) 'Autonomy and the Ethics of Biological Behaviour Modification, in Akabayashi, A. (ed.) *The Future of Bioethics: International Dialogues*, Oxford: Oxford University Press: 91–112.

Schaefer, G.O. (2011) 'What Is the Goal of Moral Engineering?' *AJOB Neuroscience* 2(4): 10–11.

Schaefer, G.O. (2015) 'Direct vs. Indirect Moral Enhancement', *Kennedy Institute of Ethics Journal* 25(3): 261-289.

Schaefer, G.O., Savulescu, J. (2019) 'Procedural Moral Enhancement', Neuroethics 12: 73–84.

Shaw, E. (2014) 'Direct Brain Interventions and Responsibility Enhancement', *Criminal Law and Philosophy* 8(1): 1-20.

Shaw, E. (2018) 'Retributivism and the Moral Enhancement of Criminals Through Brain Interventions', *Royal Institute of Philosophy Supplement*, 83: 251-270.

Shook, J.R. (2012) 'Neuroethics and the Possible Types of Moral Enhancement', *AJOB Neuroscience* 3(4): 3–14.

Sorensen, K. (2014) 'Moral Enhancement and Self-Subversion Objections', *Neuroethics* 7: 275–286.

Sparrow, R. (2014a) 'Egalitarianism and Moral Bioenhancement', *American Journal of Bioethics*, 14(4): 20–28.

Sparrow, R. (2014b) 'Better Living Through Chemistry? A Reply to Savulescu and Persson on "Moral Enhancement"', *Journal of Applied Philosophy* 31(1): 23–32.

Sparrow, R. (2014c) '(Im)Moral Technology? Thought Experiments and the Future of "Mind Control"', in Akabayashi, A. (ed.) *Towards Bioethics in 2050: International Dialogues*, Oxford: Oxford University Press.

Specker, J., Schermer, M.H.N. & Reiner, P.B. (2017) 'Public Attitudes Towards Moral Enhancement. Evidence that Means Matter Morally', *Neuroethics* 10: 405–417.

Spence, S. A. (2008) 'Can Pharmacology Help Enhance Human Morality?'*British Journal of Psychiatry* 193: 179–180.

Thaler, R.H., Sunstein, C.S. (2008) *Nudge: Improving Decisions about Wealth, Health, and Happiness*, Yale University Press.

Torres, P. (2017) Moral bioenhancement and agential risks: Good and bad outcomes. *Bioethics* 31: 691– 696.

Walker, M. (2009) 'Enhancing Genetic Virtue: A Project for Twenty-first Century Humanity? *Politics and the Life Sciences* 28(2): 27–47.

Wasserman D (2014) 'When Bad People Do Good Things: Will Moral Enhancement Make the World a Better Place? *Journal of Medical Ethics* 40:374-375.