Would Moral Enhancement Limit Freedom?

Antonio Diéguez, [1,*]

Email dieguez@uma.es

Carissa Véliz, [2]

Email carissa.veliz@philosophy.ox.ac.uk

[1] Departamento de Filosofía, Facultad de Filosofía y Letras, Universidad de Málaga, Campus de Teatinos S/N, 29071 Málaga, Spain

[2] Uehiro Centre for Practical Ethics, Faculty of Philosophy, University of Oxford, Oxford, UK

Abstract

The proposal of moral enhancement as a valuable means to face the environmental, technological and social challenges that threaten the future of humanity has been criticized by a number of authors. One of the main criticisms has been that moral enhancement would diminish our freedom. It has been said that moral enhancement would lead enhanced people to lose their 'freedom to fall', that is, it would prevent them from being able to decide to carry out some morally bad actions, and the possibility to desire and carry out these bad actions is an essential ingredient of free will, which would thus be limited or destroyed—or so the argument goes. In this paper we offer an answer to this criticism. We contend that a morally enhanced agent could lose (to a large extent) the 'freedom to fall' without losing her freedom for two reasons. First, because we do not consider that a morally well-educated person, for whom the 'freedom to fall' is a remote option, is less free than an evildoer, and there is no reason to suppose that bioenhancement introduces a significant difference here. Second, because richness in the amount of alternative possibilities of action may be restored if the stated loss is compensated with an improvement in sensitivity and lucidity that can lead to seeing new options and nuances in the remaining possible actions.

Keywords

Moral enhancement
Free will
Freedom to fall
Human bioenhancement

1. Introduction

Two seminal papers, one by Ingmar Persson and Julian Savulescu and the other by Thomas Douglas, defended the usefulness and even necessity of moral bioenhancement in order to face the challenges issued by the quick development of nuclear, biological and other technologies liable to be used as weapons of mass destruction (Persson and Savulescu 2008 and Douglas 2008). Persson and Savulescu expounded in a more detailed form their proposal in *Unfit for the Future. The Need for Moral Enhancement* (Persson and Savulescu 2012). In this book, the reasons supporting their main argument changed somewhat and other new reasons were introduced, but the essential thesis remained: The huge problems that threaten the future of humanity (environmental, political, economic and technological challenges) are likely too complex and urgent to be confronted merely by a cultural and educational change in our present attitudes and values. They would, then, require that we seriously consider the possibility of morally bioenhancing human beings. Although the precise meaning of 'moral bioenhancement' is controversial, we will assume here as an acceptable characterization that it refers to the effective improvement of an individual's moral capacities, attitudes, motivations and judgments to an above normal level through biomedical means (drugs, direct neural manipulation, or, eventually, when technological resources and skills become safe and reliable, by means of genetic engineering).[1]

AQ1

Persson and Savulescu's proposal received a number of criticisms ranging from its unnaturalness to its impracticability, its ineffectiveness, its inadequate individualistic approach, its riskiness, and its supposed immorality (cf. Agar 2013, 2015; Sparrow 2014; Beauchamp 2015; Beck 2015; Casal 2015; De Melo-Martin and Salles 2015; Hauskeller 2015; Tokens 2015, and for a review; Specker et al. 2014). Here we focus only on one of these criticisms, formulated by John Harris (2011). According to Harris, moral

bioenhancement would diminish our freedom. Moral bioenhancement would prevent human beings from carrying out morally bad actions, and although this might constitute a benefit for humankind, the possibility to desire and carry out these bad actions is an essential ingredient of free will, he argues, which would thus be limited or destroyed. Harris thinks that the freedom (to choose) to fall—to use John Milton's expression—is needed for any freedom deserving of that name:

> Autonomy surely requires not only the possibility of falling but the freedom to choose to fall, and that same autonomy gives us self-sufficiency; 'sufficient to have stood though free to fall.' […] Without the freedom to fall, good cannot be a choice; and freedom disappears and along with it virtue. There is no virtue in doing what you must. (Harris 2011, pp. 103–104).[2]

We interpret Harris' criticisms as referring to 'freedom of the will' or free will, not to 'freedom of action', since the latter notion implies the absence of an external obstacle to carry out a certain action, and it is not clear how moral enhancement could amount to such a kind of external obstacle.[3] Accordingly, Harris' criticism states that moral enhancement limits our moral autonomy insofar as our will cannot choose to carry out some candidates for possible actions; in particular it cannot choose those possibilities that imply a moral fall, i.e., the realization of (very) immoral actions. For Harris, when some morally enhanced individual does the right thing in a moral context, she could have not in fact acted otherwise, and therefore is not free.

This criticism has been answered in a number of papers: Savulescu and Persson (2012), Persson and Savulescu (2013), Douglas (2013), and DeGrazia (2014), among others. Their responses agree in a central point, namely: A proper concept of free will does not have to include the freedom to act otherwise, also known as 'the principle of alternative possibilities'. In other words, the freedom to fall is not a necessary condition of free will or of moral responsibility. We can be free even when we have no alternative possibilities to choose from. This is a thesis that was initially supported by Harry Frankfurt with a classical thought experiment (Frankfurt 1969), not very dissimilar to the later 'God-machine' thought experiment by Savulescu and Persson (2012). In Frankfurt's example, a man called Jones decides to

shoot Smith. Black learns of his plan and approves of it: he wants Jones to shoot Smith. However, Black fears that Jones might have reservations and might not go through with it. To avoid such a possibility, Black arranges things so that he will be able to secretly manipulate Jones into shooting Smith just in case Jones shows hesitation. As it happens, Jones does shoot Smith of his own accord. This thought experiment is meant to show that an agent can be morally responsible and enjoy free will even if he could not have acted otherwise. In Savulescu and Perssons' thought experiment, an intelligent computer (the God-machine) allows people to be free but is there to intervene only to prevent great harm or injustice.

In this article we give a different kind of response to Harris' challenge. We aim to show that moral enhancement would not necessarily reduce the autonomy of individuals and would not necessarily reduce the available alternatives of action to be chosen from. Firstly, we argue that a morally well-educated person, for whom the 'freedom to fall' is a remote option, is not—for this sole reason—less free than an evildoer, and there is no convincing reason to suppose that bioenhancement introduces a significant difference here. Secondly, we contend that richness in the amount of alternative possibilities of action may be restored if the adduced loss is compensated with an improvement in sensitivity and lucidity that can lead to seeing new options and nuances in the remaining possible actions.

We should warn that it is not our goal here to discuss the complex metaphysical and epistemological problem of whether free will is a real fact of the human condition or, on the contrary, a mere illusion created by our brain, as some psychologists and neuroscientists claim (Wegner 2002, and for an opposite view; Fuster 2013). We will not discuss either the relative merits of compatibilism or incompatibilism in their different modalities. These topics are enormous and we cannot hope to do full justice to them, but we think that a previous discussion of them is not essential in order to answer Harris' criticism, and he seems to think the same (Harris 2014, p. 372). Thus, for the sake of argument we will assume here that human beings have free will, or, to put it in a less metaphysically charged way, that they are autonomous agents. Just like we usually take for granted that there is free will for the purposes of daily life and legal systems, we will assume here that human beings can exercise some kind of control over their actions, such that when an individual acts freely, her actions will be the result of her decisions, instead of being the consequences of external forces.

2. Freedom to Fall and Autonomy to A~~CT~~ct

It can be argued that, even if moral enhancement worked perfectly well, and people subjected to it were transformed into moral angels, the possibility of falling would not completely disappear. Only a moral enhancement that would transform us into moral automata would completely destroy our capacity to choose, our free will. Harris admits as much. He does not claim that moral enhancement 'entirely eliminates freedom', but only that it 'would eliminate some significant measure of our freedom' (Harris 2014, p. 372).

Elsewhere, Harris makes another important concession:

> [Persson and Savulescu] and I are discussing precisely whether or not there will be some forms of moral bioenhancement that will leave room for freedom. I am sure there can be such, I have argued however that many of the forms currently being canvassed as promising do not in fact augur well for the survival either of liberty or indeed of rational strategies for seeing that good triumphs 'all things considered'. (Harris 2013, pp. 287–288).

This admission weakens the strength of Harris' initial objection. It is unclear in this context which forms of moral enhancement would eliminate some 'significant measure' of freedom.

Harris' major concern with regards to limiting free will seems to be focused on moral enhancement directed toward the improvement of motivation. The main reason that Harris adduces against enhancing motivation is the following:

> My claim is that while the influence is indeed on motivation (if that term is understood, as it sometimes is, simply as a mainspring of action), this 'motivation' does not meet standards of moral reasoning for the simple and sufficient reason that it does not meet standards of reasoning at all. The intervention is designed to bypass reasoning and act directly on attitudes. When such attitudes are manipulated, not only is freedom subverted but also morality is bypassed. (Harris

> 2014, p. 372).

He then adds:

> My quarrel is with the process whereby those apparent values and preferences are acquired and become stable, and my purpose is to try to enable those who acquire them to be free to continue to test them by seeing whether holding them does indeed conduce to the good. […] The freedom of which I (and I believe Milton) speak, the freedom to fall is the freedom to decide whether or not to fall for reasons, which have to do with what is best 'all things considered'. Anything, which by influencing attitudes or emotional responses inhibits that ability significantly, is inimical to freedom. (Harris 2014, p. 373).

Harris specifies, then, that not every moral enhancement limits freedom. However, most cases mentioned as possible ways of moral enhancement (drugs that reduce aggressiveness or xenophobia, for instance) do limit it, he thinks, because they 'bypass' moral reasoning and act directly on attitudes and motivations.

It seems clear that Harris changes the focus of his objection in his later texts. In order to clarify this change and evaluate its merits, it is worth keeping in mind an important distinction. Although debatable, it has been common among experts to characterize free will as having at least two central features (Ekstrom 2012):

i. self-determination, or autonomy; and

ii. the availability of alternative possibilities.

The first feature points out that we tend to consider free actions only those that depend on the will of the agent, and are therefore governed by the agent and not by an external will. The second feature refers to the possibility to act otherwise, although, as we mentioned previously, disagreements on this point remain.[4] We interpret that Harris' initial criticism—that moral enhancement

impedes the *choice* to fall—is concerned directly with the infringement of the second clause for freedom of agency, and only indirectly with the first. In contrast, his later criticism—that moral enhancement bypasses reasoning—is concerned with the infringement of the first clause. As we will see, however, some significant forms of moral enhancement would infringe neither of these two requirements.

3. Is it Possible to be Morally Enhanced and Free?

We will first discuss Harris' criticism related to the lack of autonomy in some conceivable cases of moral enhancement.

Let's take as an illustration a recent experiment that can be considered a case of moral enhancement insofar as the role of biased racial stereotypes of the experimental subjects were modulated by technological means (Sellaro et al. 2015). More precisely, this experiment shows that a certain kind of transcranial current stimulation of medial prefrontal cortex (mPFC)—an area linked to the processing of socio-cognitive information, and among other things "implicated in the representation of an individual's traits, preferences and mental states during the formation of impression about other people" (p. 2)—decreases implicit biased attitudes toward out-group members.[5] The experimenters found that this kind of electric stimulation decreased the racial biases of 60 Dutch students with regards to the names of Moroccan persons. The authors suggest that the medial prefrontal cortex may be critical in counteracting stereotypes activation and its stimulation can initiate cognitive-control processes aimed to override unwanted responses driven by stereotypes activation. According to them, their results "are consistent with previous findings showing that increasing cognitive control may overcome negative bias toward members of social out-groups".

If this conclusion is right, then this experiment is a good example of how moral enhancement could produce more free and cognitive-controlled behavior. Through brain stimulation, impulsive and unconscious biased attitudes are temporarily replaced by rationally controlled reactions. In other words, this experiment shows that it is possible to carry out moral enhancement by procedures able to increase self-control, reason-responsive behavior, and personal autonomy. It is not that the implicit biases are rationalized and suppressed. These biases remain working unconsciously, but their influence is weakened since it is counteracted by a better cognitive control of individuals' responses. They respond taking into account more

clearly the genuine reasons related to the situation, without having their capacities distorted by these biases.

There is a tendency to see moral enhancement as a way to change by means of technology the (bad) desires or motivations of an individual for other (good) desires or motivations, where this looks like the product of external manipulation. But it is perfectly possible that most accepted varieties of moral enhancement in the future would be directed to the improvement of our reflective capacities, so that we do not act so frequently according to impulses, or following comminatory aversions, or strong wishes, but rather as a result of decisions elaborated after careful deliberation. As Shaefer, Kahane and Savulescu (2014), among others, have argued, these reasoning, deliberative and evaluative capacities are closely related to the autonomy of individuals. As long as future intended moral enhancement seeks this kind of brain modifications, it does not seem that it should induce fear of a diminishment of freedom.

It could be adduced that this empirical study is not helpful for our purposes because Harris was primarily concerned about modulating emotions, whereas the study is talking about enhancing cognitive control, and therefore something Harris might, in principle, actually be in favor of. It is true that Harris might accept that enhancing moral cognitive control over prejudices would not limit our freedom. But in fact this experiment shows quite a different thing. It shows that an external influence (an electrical stimulation of one part of brain) can make more difficult some morally relevant mental representations and emotional reactions. The increase of cognitive control is not produced through a direct improving of our explicit and conscious “reasons” to act, or of our sensitivity to appreciate better moral arguments, or blocking the formation of biased attitudes, but weakening the power of these prejudiced attitudes already unconsciously present in the subjects. The authors acknowledge that the precise mechanisms are not well known, but the final effect is to overcome unwanted biased responses. Taking into account all the information that the authors provide in the article, we understand this process as an improvement of the control of our emotional responses which in fact bypasses reasoning. Put simply, after the stimulation the subject can avoid some reactions which are somewhat undesirable for her, regardless of whether she has reasoned about the real cogency of the motives for these reactions or not. Reactions driven by stereotypes and prejudices are weakened as a result of the stimulation, and not through careful reflection, by providing

subjects with persuasive arguments, or making them embarrassedly aware of their implicit biases. The brain stimulation changes the subject's previous xenophobic attitudes, the control over her biased behavior, but does not change the underlying beliefs that she may have. Obviously, one study, on its own, cannot provide a definitive refutation of Harris' theses, but it is enough to show that there can be some possible forms of moral enhancement that in spite of bypassing reason, do not diminish freedom.

It must be admitted that it is intuitive to think that, in the case of virtuous people who do the right thing, the ultimate causes of their moral actions seem to be constitutive of the agent in some deeper way than in cases of morally enhanced agents. But the question is if this intuitive difference warrants the conclusion that the morally enhanced agent is not free because in some sense she has psychological states or dispositions that are not genuinely felt and assumed as her own. We think that this empirical study provides good reasons to claim that this is not the case. Between the extremes of the morally virtuous agent who has cultivated her personality throughout her life and the "robotic" agent whose actions are the direct result of external manipulations or decisions, there is much room for freedom.

As for Harris' initial concern about the necessity to keep the 'freedom to fall' in order to preserve our freedom, any attempt to find a response has to be a little more complex. First of all, we agree with Persson and Savulescu (2013, p. 128) that "[p]eople who are morally good and always try to do what they regard as right are not necessarily less free than those who sometimes fail to do so". Particularly, a morally good person is not necessarily less free than an immoral person despite the fact that, in some sense, it could be said that this person is not free to fall, since although the possibility to choose a very immoral course of action is always open for her, her cultivated character, her aversion to evil and her strong will, make her not fall. Any normal moral agent feels compelled by a set of moral principles and values. These principles and values constitute and characterize her as an individual with a certain culture, personality, and life project. More than a limitation to freedom, these principles and values are the basis and guide to genuine free action.

Likewise, there seems to be no good reason why a morally enhanced person, who would not have the 'freedom to fall' in this concrete sense is not free at all, or less free than a virtuous person who is not morally enhanced. For Aristotle (1999 NE VII), for example, the virtuous is the person who knows

what the right thing to do is, does it, and does it without conflict. For the truly virtuous temptations do not exist. In this sense, the virtuous is the person who decides that she will not even consider evil options as possible courses of action. But this does not mean she is not free. On the contrary, she is freer because she has perfect autonomy.

There are other examples where (non-biomedical) moral enhancement does not seem to lead to less freedom. The rate of homicides has notably decreased in Europe over the last centuries as a consequence of a number of disparate causes: the improvement and extension of the educational system, the higher level of life and welfare, the introduction of legal, institutional and political reforms, the increase of equality, and so on. All of these can be considered cultural innovations that contributed to moral enhancement and to reducing our freedom to fall.[6] Nevertheless, it would be preposterous to consider that these innovations cause a loss in the freedom of people just because a very immoral act (homicide) was practiced by a smaller percentage of people after their implementation. Presumably, even though homicide is nowadays a more remote possibility for the average person, people cannot be seen as morally less autonomous compared with their equivalents in previous centuries. Having better external circumstances make homicide less of an attractive option. In Aristotle's case, the virtuous person has banished evil from her range of actions through self-discipline and habituation. In both cases, free will seems to remain intact.

Consider the following case. Jane has one dollar and she goes to spend it in a one dollar shop. In World 1, the shop Jane goes to has ninety-nine items to choose from. Jane chooses chocolate *x*. In World 2, the shop Jane goes to has a hundred items to choose from. The extra item is a bag of rat poison. Jane does not have a rat problem in her house and has no use for rat poison. Jane chooses chocolate x. It would be absurd to claim that Jane is less free in World 1 merely because she cannot choose a useless and toxic bag of rat poison. Having less options does not necessarily imply less free will. As long as we have enough acceptable options to choose from, our freedom remains untouched. Likewise, even if 'falling' is not an option in some scenarios for a morally enhanced person because this person now perceives that she has no need to choose this option, that the final result would be useless or detrimental, this person is no less free.

It is interesting to note that, in fact, our last point can be easily inferred if we accept the influential view that freedom consists essentially in the capacity to

respond to reasons (Fischer and Ravizza 1998, Chaps. 2 and 3, and for a discussion, see; McKenna 2012 and McKenna and Pereboom 2016).

Furthermore, it is not even clear that moral enhancement would actually lead to having fewer options available. Let us imagine a situation in which a moral agent can act without any external coercion. Let us also assume that, in spite of Frankfurt's thought experiment, free agency depends on the fact that the agent be able to choose between several alternatives. It is true, by definition, that a morally enhanced agent would hardly choose at least some (very) immoral possible courses of action. To put it more controversially: in extreme cases a morally enhanced person will not be able to choose the most evil course of action. However, this does not imply that the number of available alternatives to be taken into consideration by the agent has necessarily diminished. A morally well-educated and sensitive person can perceive nuances that other people cannot appreciate, and this capacity may enrich the set of alternative possibilities of actions available to her.

This person will surely discard a number of immoral actions that the morally insensitive person probably accepts as practicable. However, her more developed and sophisticated sensitivity for moral nuances might permit her to realize that the remaining possibilities of action are more diverse and numerous than it could be thought if given less attentive consideration. In a fractal image, it is possible to lose a part of the whole structure while keeping all its complexity in a more detailed view of the remaining parts. In the same way, even if we accept that a morally enhanced person has lost the 'freedom to fall', it is reasonable to think that her improved feelings, reasoning, and motivations, contribute to maintaining and even increasing the complexity and diversity of the spectrum of morally relevant alternative possibilities of action. These possibilities are not only *perceived* possibilities—they become *actual* possibilities as soon as the agent comes to think of them as possibilities of action. As the neuroscientist Joaquín Fuster claims, "[b]oth increased inputs and increased outputs add up to more freedom for the organism". He goes on to clarify in a later passage: "[a] better understanding and expression of our thoughts opens the range of our options. The key lies in the power of language to create new affordances, new possibilities of action […]"(Fuster 2013, pp. 11–169). As a simple illustration, a morally unenhanced person might see only two options when faced with a particularly strong conflict with a rival who is making her life impossible: either leave her city, and cede victory to her enemy, or kill her enemy. After getting morally enhanced—thus

overcoming emotions like hate and fear, and gaining a better understanding of the causes for the conflict—she might be incapable of choosing the option of killing her opponent, but may suddenly realize there is an array of possible options she could choose (e.g. persuade her enemy, make amends, etc.).

Thus far, we have argued that moral enhancement need not amount to either being less free or to having fewer options. One might still think, however, that something like freedom of mind is jeopardized with moral enhancement. More recently, a similar objection to that of Harris has been set out against moral enhancement. Responding to Savulescu and Persson's 2012 article, Christoph Bublitz has pointed out that there are some cases of possible decrease of mental freedom that these authors do not contemplate. As he explains:

> [Savulescu and Persson] too easily pass over a very influential variation of compatibilism which denies responsibility if persons act on preferences that were brought about through manipulation. This position, most notably championed by John Fischer and Mark Ravizza, has emerged from controversies around 'nefarious neurosurgeons' who manipulate an agent's preferences so that she commits evil deeds. Manipulated agents, so the reasoning goes, should not be held responsible for their actions despite the fact that they satisfy all usually cited conditions of responsibility. Fischer & Ravizza's central idea is that autonomy is a historical phenomenon, which means that apart from sufficient capacities in the moment of action, the genesis of preferences leading to action is relevant. Persons may acquire preferences on autonomy undermining routes, especially those that bypass capacities for reason. Therefore, two persons may possess identical psychological structures in a given moment, yet their responsibility might be assessed differently. What matters is how they acquired their preferences and became the persons they are. (Bublitz 2016, p. 94).

The main objection posed by Bublitz is that, even if moral bioenhancement did not imply a diminishing of free will, it could infringe some senses of 'mental freedom'. He defines freedom of mind as "the freedom of a person to

use her mental capacities as she pleases, free from external interferences and internal impediments" (2016, p. 94). In order to clarify how moral bioenhancement could affect this mental freedom, he puts the following example:

> [S]uppose someone spikes your drink with a substance that mildly impairs your cognitive capacities, not dramatically, only of the sort that everyone experiences on a bad day; or that induces distracting sensations while you seek to concentrate. These newly induced mental states are not of a kind that undermines free will or resolves you from responsibility. Still, you may legitimately complain about some detrimental effect to your mind, and normatively, about the infringement with some protected interest. The interest in question is not free will — but freedom of mind. (Bublitz 2015, p. 95)

Bublitz seems to reiterate Harris' second concern, already analyzed here. Our response, thus, is the same. Admittedly, the way in which morally enhanced persons acquire their preferences matters, and it matters a great deal. The historical path that the agent follows in order to make a decision is relevant to decide on her freedom, as Fischer and Ravizza (1998) have argued.[7] This is particularly clear when we consider the very first action the agent performs immediately after an unconsented manipulation. Fischer and Ravizza, and other historical compatibilists, have emphasized that an agent cannot be free (nor morally responsible) if she has not had time enough to reflect on her new situation, her new capacities and worldview. In this case, the agent suffers damage in her evaluative and volitional capacities. If subjects do not accept moral enhancement voluntarily or if this enhancement is as manipulative as to "bypass capacities for reason", then it could limit or nullify the autonomy of the agent. It cannot be denied that increasing or diminishing the empathy of a person by means of hypnosis or of alcohol should be considered as a limitation of freedom, since her actions would then obey the will of the hypnotizer or would be the effect of a drug weakening the voluntary control of actions. However, as we have argued, there are other conceivable modalities of moral enhancement that could increase autonomy and would affect not only the very first action after the enhancement, but many others in the future about which agents would have time to think carefully. It is important to pay attention to the potential threat of mental manipulation,

particularly when the manipulation is imposed without the consent of the subject by an external agent. Manipulative techniques, however, are not the only possible way to achieve moral bioenhancement. Some techniques could be used to improve the ability to respond to reasons, both short and long term. The agent would then become more 'reason-responsive', not less. This is something that Bublitz admits, since in response to Thomas Douglas he concedes that an intervention reducing aggressive impulses "supposedly has beneficial effects on mental control". Moreover, it would be excessive to require complete mental freedom in order to consider an action as genuinely free. Complete freedom "from external interferences and internal impediments" is likely a requirement impossible to fulfill, since it would be enough that we take into account in our deliberations the authoritative opinion (not the arguments) of a prestigious person who disagrees with us in order for us not to be genuinely free in this sense. If, as Bublitz says, one aspect of mental freedom is conscious control over one's mind, this control is arguably always a matter of degrees.

4. Conclusion

Assuming, for the sake of the argument, that human beings are free agents (i.e. that freedom is not an illusion), the implementation of moral bioenhancement would not necessarily imply a reduction in freedom. Although there might be some possible forms of moral bioenhancement which transform individuals into mere automata, as some critics have emphasized, there are other conceivable forms that would not nullify the basic conditions for a free agency: they would not impede the agent's initiative, they would not make the actions the necessary consequence of the enhancement, and they would not turn actions into something completely predictable. Moral bioenhancement could be aimed to act upon the level of deliberation and assessment, and not directly upon the level of decision. In that case, bioenhancement might improve the cognitive control of the process of decision making, and therefore it would increase the autonomy of individuals. But even if in some cases bioenhancement could bypass rational deliberation and act upon motivations or feelings, this would not necessarily lead to a less free agency. Suppose that the enhancement in question is limited to adding a negative evaluative burden to some possible actions previously disapproved by the agent. As a result of it, the agent loses motivation for the realization of these actions, but this effect is not very dissimilar to that produced by moral education through reading or watching documentaries—procedures that

nobody considers as limiting of freedom.

On the other hand, a morally enhanced agent could lose to a large extent the 'freedom to fall'—in the sense that she would manifest a strong tendency to refuse the most immoral alternatives of action—without losing her freedom, given that: (a) someone who acts morally virtuously need not be less free than a wrongdoer, and there is no reason to suppose that moral bioenhancement would introduce a significant difference, and (b) increasing moral capacities through moral bioenhancement may increase moral affordances, in the sense that the agent may perceive new possibilities of action which would make up for the loss of the worst alternatives.

Acknowledgements

We are very grateful to an anonymous referee for very helpful comments.Research for this work has been supported by the research projects FFI2012-37354 (Spanish Government), HUM-0264 and HUM-7248 (Junta de Andalucía), and project PAPIIT-UNAM IN403613 (Mexico).

References

Agar N (2013) Why is it possible to enhance moral status and why doing so is wrong? J Med Ethics 39:67–74

Agar N (2015) Moral bioenhancement is dangerous. J Med Ethics 41(4):343–345

Aristotle (1999) Nicomachean Ethics. Terence Irwin (trans.). Indianapolis: Hackett

Beauchamp TL (2015) Are we unfit for the future? J Med Ethics 41:346–348

Beck B (2015) Conceptual and practical problems of moral enhancement. Bioethics 29(4):233–240

Bublitz C (2016) Moral enhancement and mental freedom. J Appl Philos 33(1):88–106

Casal P (2015) On not taking men as they are: reflections on moral bioenhancement. J Med Ethics 41:340–342

De Melo-Martin I, Salles A (2015) Moral bioenhancement: much ado about nothing? Bioethics 29(4): 223–232

DeGrazia D (2014) Moral enhancement, freedom, and what we (should) value in moral behavior. J Med Ethics 40:361–368

Douglas T (2008) Moral enhancement. J Appl Philos 25(3):228–245

Douglas T (2013) Moral enhancement via direct emotion modulation: a reply to John Harris. Bioethics 27(3):160–168

Ekstrom LW (2012) Free will is not a mystery. In: Kane R (ed) The oxford handbook of free will: second edition. Oxford University Press, Oxford, pp 366–380

Fischer JM (2006) My way: essays on moral responsibility. Oxford University Press, Oxford

Fischer JM, Ravizza M (1998) Responsibility and control: a theory of moral responsibility. Cambridge University Press, Cambridge

Frankfurt H (1969) Alternate possibilities and moral responsibility. J Philos 66:829–839

Fuster JM (2013) The neuroscience of freedom and creativity. Cambridge University Press, Cambridge

Harris J (2011) Moral enhancement and freedom. Bioethics 25(2):102–111

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

Harris J (2014) Taking liberties with free fall. J Med Ethics 40(6):371–374

Hauskeller M (2015) Being good enough to prevent the worst. J Med Ethics 41(4):289–290

Levy N (2017) Implicit bias and moral responsibility: probing the data. Research 94(1):3–26

McKenna M (2012) Contemporary compatibilism: mesh theories and reasons responsive theories. In: Kane R (ed) The oxford handbook of free will: second edition. Oxford University Press, Oxford, pp 175–198

McKenna M, Pereboom D (2016) Free will. A contemporary introduction. Routledge, New York

Persson I, Savulescu J (2008) The perils of cognitive enhancement and the urgent imperative to enhance the moral character of humanity. J Appl Philos 25(3):162–177

Persson I, Savulescu J (2012) Unfit for the future. The need for moral enhancement. Oxford University Press, Oxford

Persson I, Savulescu J (2013) Getting moral enhancement right: the desirability of moral bioenhancement. Bioethics 27(3):124–131

Raus K et al (2014) On defining moral enhancement: a clarificatory taxonomy. Neuroethics 7(3): 263–273

Savulescu J, Persson I (2012) Moral enhancement, freedom and the god machine. Monist 95(3):399–421

Sellaro R et al (2015) Reducing prejudice through brain stimulation. Brain Stimul 8:891–897

Shaefer GO, Kahane G, Savulescu J (2014) Autonomy and Enhancement. Neuroethics 7:123–136

Sparrow R (2014) Better living through chemistry? A reply to Savulescu and Persson on ‘moral enhancement’. J Appl Philos 31(2):23–32

Specker J et al (2014) The ethical desirability of moral bioenhancement: a review of reasons. BMC Med Ethics 15(67):1–17

Tokens R (2015) ‘My Child will never initiate ultimate harm’: an aargument against moral enhancement. J Med Ethics 41:245–251

Wegner DM (2002) The illusion of conscious will. The MIT Press, Cambridge

[1] For a clarification of the different meanings of ‘moral enhancement’, see Raus et al. (2014). We assume here that there is no precise way to determine, for any possible case, whether or not it could be included under this definition. As these authors rightly point out, “as there is no objective way of determining what falls within the range of ‘normal’ moral behaviour or functioning of moral capacities, every choice of cut-off point is, necessarily, a normative one. Even the question as to what constitutes ‘moral capacities’ has no straightforward answer.” (p. 269).

[2] Harris’ position changed in later papers, ~~and afterwards Harris seems to have himself retreated into silence on the issue,~~ but we think that this argument is interesting enough as to deserve an answer regardless of Harris`s final opinion.

[3] Presumably, moral bioenhancement would not work through impeding physical actions of certain kinds. Moral bioenhancement must modify our inclinations to do good, rather than our capabilities to move.

[4] For a review of the main criticisms, see Fischer (2006, Chap. 2), and McKenna and Pereboom (2016, Chap. 5).

[5] For an updated discussion of the character of implicit biases and their effects on behavior and moral responsibility, see Levy (2017).

[6] Obviously, these social and institutional innovations could have also contributed to human cognitive enhancement, and in this sense, some of them do not bypass reasoning, but our point still stands. With this example we mean to illustrate that the mere fact of diminishing our freedom to fall does not diminish our free will, which was Harris’ initial concern.

[7] This is, however, a controversial issue. For a discussion, see McKenna and Pereboom (2016, Chap. 8).