Compulsory Moral Bioenhancement Should Be Covert

ABSTRACT: Some theorists argue that moral bioenhancement ought to be compulsory. I take this argument one step further, arguing that *if* moral bioenhancement ought to be compulsory, then its administration ought to be covert rather than overt. This is to say that it is morally preferable for compulsory moral bioenhancement to be administered without the recipients knowing that they are receiving the enhancement. My argument for this is that if moral bioenhancement ought to be compulsory, then its administration is a matter of public health, and for this reason should be governed by public health ethics. I argue that the covert administration of a compulsory moral bioenhancement program better conforms to public health ethics than does an overt compulsory program. In particular, a covert compulsory program promotes values such as liberty, utility, equality, and autonomy better than an overt program does. Thus, a covert compulsory moral bioenhancement program is morally preferable to an overt moral bioenhancement program.

Keywords: moral enhancement; public health ethics; autonomy; public policy; harm

Advancements in technology and globalization create many opportunities. Among them are not only greater opportunities to enhance our capacities for cognition, but also to use those enhanced capacities to cause "ultimate harm," which is an event or series of events that brings about either the annihilation of humanity or a condition of living that is so poor that life is not worthwhile.[1] Global in scope, ultimate harm is characterized by high rates of death, and, for those that survive, immense suffering. Examples of potential catalysts of ultimate harm are familiar and include global warming, terrorists, and short-sighted scientists unleashing a plague. Some argue that as cognitive enhancement increases in scope and availability, the potential for ultimate harm is greater, especially when conjoined with some of the possible catalysts of ultimate harm.[2] But an event or series of events doesn't have to be caused by an agent or agents to be ultimate harm. A large meteor or supernova could also bring about ultimate harm.

[1] 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

[2] Op. cit. note 1.

Persson and Savulescu claim that in order to prevent the greater possibility of ultimate harm, it is necessary for humans to be morally enhanced.[3] More specifically, it is necessary to morally bioenhance the population in order to prevent ultimate harm. Moral bioenhancement is the potential practice of influencing a person's moral behavior by way of biological intervention upon their moral attitudes, motivations, or dispositions. The technology that may permit moral bioenhancement is on the scale between non-existent and nascent, but common examples of potential interventions include infusing water supplies with pharmaceuticals that enhance empathy or altruism or otherwise intervening on a person's emotions or motivations, in an attempt to influence the person's moral behavior.

The proposal of engaging in moral bioenhancement has since generated a lot of controversy, resulting in a substantial literature debating the ethics of moral bioenhancement. Among the ethical issues with moral bioenhancement are its impact on autonomy, whether moral bioenhancement will result in a more or less egalitarian society, whether it will erode personal identity, and whether it will benefit others or just the people who are morally enhanced.[4] There has also been debate around the empirical issues of whether a means of moral bioenhancement is or ever will be feasible and whether it really is necessary.[5] Others have argued that if moral bioenhancement is

[3] Persson & Savulescu op cit.; Persson, I., & Savulescu, J. (2012). *Unfit for the Future: The Need for Moral Enhancement*. Oxford University Press. Persson, I., & Savulescu, J. (2013). Should moral bioenhancement be compulsory? Reply to Vojin Rakic. *Journal of Medical Ethics*.

[4] Douglas, T. (2008). Moral enhancement. *Journal of Applied Philosophy*, *25*(3), 228–245; Sparrow, R. (2014). Egalitarianism and moral bioenhancement. *The American Journal of Bioethics : AJOB*, *14*(4), 20–8;

[5] Crockett, M. J. (2014). Moral bioenhancement: a neuroscientific perspective. *Journal of Medical Ethics*, *40*(6), 370–1; Harris, J. (2011). Moral enhancement and freedom. *Bioethics*, *25*(2), 102–111; Sorensen, K. (2014). Moral enhancement and self-subversion objections. *Neuroethics*, *7*(3), 275–286.; Wiseman, H. (2014a). SSRIs as Moral Enhancement Interventions: A Practical Dead End. *AJOB Neuroscience*, *5*(3), 21–30.; Wiseman, H. (2016). *The Myth of the Moral Brain: The Limits of Moral Enhancement*. MIT Press.

necessary to prevent ultimate harm, then moral bioenhancement should be compulsory.[6] I adopt this argument, but go a step further. Not only should moral bioenhancement be compulsory, it should also be covert, conducted without the knowledge of those who are being enhanced. I have previously argued on the grounds that it is the only way for a program to be effective for it to be covert (Author, redacted). Here I argue on ethical grounds that a moral bioenhancement program ought to be covert, if it ought to be compulsory.

My argument rests on two assumptions. The first is that moral bioenhancement is necessary to prevent ultimate harm. Whether moral bioenhancement is so necessary is an empirical matter. Here, I assume it is necessary. The second assumption is that whatever moral bioenhancement is used will be safe and effective. Moral bioenhancement will be safe and effective when it can reliably intervene on a person's moral capacities and do so without causing serious adverse events. Whether this assumption is true is also an empirical matter.

Admittedly, these are demanding assumptions, and for some readers this demand may push what follows into the merely theoretical. Perhaps I give too much credit to scientists and too little to the moral motivations of large groups of people, but I don't find the assumptions to be merely theoretical. Moreover, most of the recent discussion of the ethics of moral bioenhancement has proceeded from the assumption that moral bioenhancement can be safe and effective. It is a common assumption, and without it we can all stop discussing the ethics of moral bioenhancement until the science is settled. However, even if the argument that follows is merely theoretical, then it at least represents the limiting case from which it is possible to reason about other potential interventions upon moral behavior.

The argument that follows is intended to support the claim that *if* moral bioenhancement ought to be compulsory, it ought to be covert. I do not offer a vigorous defense of the antecedent of this conditional, that moral bioenhancement ought to be compulsory. Others have defended this elsewhere, but I do review the argument for this

[6] Persson & Savulescu (2008) op. cit. note 1

position in the first section. Second, I argue that contrary to the focus of the literature on moral bioenhancement, moral bioenhancement is a matter of public health and is therefore subject to ethical norms that guide decisions on public health interventions rather than the ethical norms that guide individual medical decisions. I then introduce the frameworks for public health ethics. The final step is to argue that according to these frameworks, moral bioenhancement ought to be covert, if it is to be compulsory. Compared to an overt program, a covert program is better. This conclusion is likely to be unacceptable to some. I aim to establish the implication of covert moral bioenhancement from compulsory moral bioenhancement. If the consequent of this implication is unacceptable, it is a short step to the notion that moral bioenhancement should be voluntary. And if it shouldn't be voluntary, the only alternative is to forego moral bioenhancement completely. I conclude with a discussion of these inferences.

1. Compulsory Moral Bioenhancement

Given that the costs of not preventing ultimate harm are indefinitely high[7], there is no intervention the costs of which would outweigh utility of the prevention of ultimate harm. Thus, if an intervention is necessary to prevent ultimate harm, and the intervention will actually prevent ultimate harm, then that intervention ought to be carried out, because the cost of not doing so is indefinitely high. Moral bioenhancement is necessary, because as cognitive enhancement makes causing ultimate harm more accessible to nefarious moral agents, ultimate harm is much more likely, unless everyone is enhanced.[8] Where it used to require an extraordinarily coordinated effort to cause ultimate harm, now, or in the near future, it only takes one person. Thus, moral bioenhancement ought to be compulsory for everyone.

[7] The costs of ultimate harm are so high because (a) not existing at all or existing in a state in which it would be better to not exist is the worst event, state, or intervention and (b) if ultimate harm were to occur it is impossible to know the goodness that such an occurrence would prevent (Persson & Savulescu op. cit. note 1).

[8] I discuss in section five the likelihood and potential inevitability of ultimate harm.

To get the conclusion that moral bioenhancement ought to be compulsory, three propositions must be true: that the costs of not preventing ultimate harm are indefinitely high; that moral bioenhancement is causally necessary to prevent ultimate harm; and that moral bioenhancement is safe and effective. That the costs of not preventing ultimate harm are indefinitely high is disputable. If ultimate harm has occurred, then most of us are either dead or in a state of living that is so bad that death would be an improvement. But that this is the worst possible outcome is contingent on valuing life and freedom from intense, prolonged suffering. If there are things more valuable than this, and ultimate harm doesn't cause their loss, then ultimate harm may not be the worst possible outcome.

For the sake of the argument that follows, I assume that moral bioenhancement is causally necessary to prevent ultimate harm, which is to say that there are no alternatives that exclude moral bioenhancement. If moral bioenhancement is unnecessary to prevent ultimate harm, it doesn't follow that it shouldn't be compulsory.[9] It may be logically possible that ultimate harm can be prevented by a combination of many programs rather than moral bioenhancement. My claim is one of causal necessity, not logical necessity.

That moral bioenhancement ought to be compulsory doesn't follow simply from the costs of not preventing ultimate harm being indefinitely high and the necessity of moral bioenhancement to prevent it. Moral bioenhancement must also be safe and effective. If moral bioenhancement is not safe and effective, then it may be that moral bioenhancement is necessary to prevent ultimate harm, though not sufficient. And if it's not sufficient, it may be that ultimate harm is inevitable, so the costs of making moral bioenhancement compulsory (whatever they happen to be) needlessly add to the costs of ultimate harm. However, if the costs of not preventing ultimate harm are indefinitely high, and if moral bioenhancement is both necessary to prevent ultimate harm *and* safe and effective, then it follows that moral bioenhancement ought to be compulsory.

---

[9] Rakić (Voluntary moral bioenhancement is a solution to Sparrow's concerns. *The American Journal of Bioethics*, *14*(4), (2014) 37–8) argues that moral bioenhancement should not be compulsory. I discuss in more detail below how the alternative of voluntary moral bioenhancement articulates with my argument that moral bioenhancement, if compulsory, should be covert.

It is clear that there is not currently any type of moral bioenhancement that is so safe and so effective that it would meet standards of safety and effectiveness. This does not mean that it never will be safe and effective. But if it is impossible for moral bioenhancement to be safe and effective, then *any* argument for the use of moral bioenhancement would be significantly weakened.

2. Compulsory Moral Bioenhancement is a Matter of Public Health

The case that a safe, effective, and necessary moral bioenhancement is a matter of public health is also rather straightforward. Moral bioenhancement is a health intervention. The main aim of any health intervention is to improve well being, but sometimes the intervention aims to improve the well being of an individual and the individual's close social circle, while other times the intervention aims to improve the well being of populations and only secondarily the well being of a particular individual.

One significant ethical difference between interventions aimed at individuals and interventions aimed at populations is the range of values that one considers in whether to administer the intervention. In the case of an intervention on an individual's health, these considerations include the patient's preferences, the promotion of what's good for the patient and the prevention of what's harmful, and the personal relationships the patient has with others. In the case of interventions on a population's health, the values of concern are those that apply to the whole population, such as whether the intervention is administered fairly, promotes liberty or opportunity, uses resources appropriately, and engenders a net gain of well being.

The aims of a compulsory moral bioenhancement program are to prevent ultimate harm. Ultimate harm is a condition of large groups of people. Compulsory moral bioenhancement is a health intervention aimed at preventing large groups of people from being harmed, an intervention aimed at the public's health. Therefore, whether a compulsory moral bioenhancement program ought to be instituted is a matter of public health, and its ethical permissibility ought to be determined by the ethical frameworks of public health interventions.

The ethics of using moral bioenhancement to prevent ultimate harm do not resemble the ethics of a terminal patient's choice to die on one's own terms, or the decision of a pregnant woman to abort a fetus, or an infertile couple's decision to use new technologies to reproduce, or even an athlete's choice to use performance-enhancing drugs. The primary difference is that the effects of moral bioenhancement are much more widespread than the effects of an individual's medical treatment. Thus, the frameworks of biomedical ethical reasoning that have been developed to investigate the ethics of decisions that are limited to the individual or the individual's close social group should not be expected to also adequately inform the ethics of issues that are of such great public concern.

The move to public health ethics is significant, because the values considered in administering the intervention are different from those of interventions aimed at individuals. In particular, for public health ethics the restriction of individual liberty is more permissible than it would be if the public's health weren't at risk. This fact is obvious when we consider the ethical permissibility of quarantining a patient with a highly infectious disease and the ethical permissibility of quarantining a patient with a disease that causes equal individual harms but is not infectious at all. In the former case the quarantine is ethically permissible; in the latter it isn't. The restriction in liberty is justified by the greater expected utility of preventing the spread of infection. Other examples of the permissibility of liberty restrictions for the sake of the population's health are the fluoridation of public water supplies and mandating the vaccination of children, if they are to attend public schools.

3. Public Health Ethics Frameworks

The development of frameworks for public health ethics is relatively recent. Several distinct frameworks have been developed, advocating for the consideration of a wide variety of values that ought to be considered when implementing a public health program. Given that public health interventions are intertwined with liberty interests and

potentially competing interests in utility, it is not surprising that the frameworks all call for the consideration of these two values.

Ross Upshur developed a framework that proceeds in the way similar to other discussions in bioethics, via principles of action.[10] He claims there are four of these: that liberty restrictions can only be justified by the prevention of anticipated harm, that the least liberty-restricting means must be used, that public health officials must reciprocate for compliance with the program, and that the program should be transparent.

Other authors formulate their frameworks as a way to balance a range of potentially competing values. Childress et al. suggest that the values that should be considered include maximizing benefits and minimizing harms (utility), distributing benefits and burdens fairly, respecting autonomous choices and liberty of action, protecting privacy, keeping promises, being transparent, and building and maintaining trust.[11] Like Upshur they also give several conditions for infringing on one of these values, or how to balance the different values. In fact, the conditions are similar in content to Upshur's principles. The conditions are that the infringement be necessary, effective in achieving the desired public health benefit, that benefits proportionally outweigh the infringement, that it be the least restrictive means, and that it be publically justified.

The frameworks of Upshur and Childress et al. place the balance of utility and liberty as the central feature of determining the ethics of a public health program. Others broaden the range of values even further. For examples, Selgelid argues that rather than adhere to the principles above, the ethics of a public health program should be determined by the balance of utility, liberty, and equality, as it is possible that sometimes the promotion of equality outweighs potential liberty restrictions (it's possible that each of the three values outweighs the other two).[12] And Grill and Dawson argue that a wider

[10] Upshur, R. E. G. (2002). Principles for the justification of public health intervention. *Canadian Journal of Public Health*, *93*(2), 101–103

[11] Childress, J. et al. (2002). Public health ethics: mapping the terrain. *The Journal of Law, Medicine & Ethics*, *30*(2).

[12] Selgelid, M. J. (2009). A moderate pluralist approach to public health policy and ethics. *Public Health Ethics*, *2*(2), 195–205.

range of values ought to be considered.[13] They write that "some values that we would ourselves include in many public health decisions are individual health, population health, health equality, individual liberty, solidarity, social trust, and material wellbeing," as well as the honoring of duties, the non-infringement of rights, and the expression of virtues. They go on to provide a method of making decisions about public health programs that is neutral between these values.

Collecting the different frameworks, we have several principles or conditions, and a wide range of values to consider. The conditions are that it's necessary, effective, proportional, publically justified, and the least restrictive means. The values considered are liberty, utility, equality, fairness, transparency, promotion of individual and population health, trust, solidarity, respectful of autonomy, and honoring of duties.

4. Covert Moral Bioenhancement

The present issue is not whether the public health program of administering moral bioenhancement ought to occur; it's a matter of how it should occur. Let us suppose that if it were to occur overtly, it would occur similarly to vaccination programs for children: at the age where the moral bioenhancement is safe and effective, children would receive the moral bioenhancement from their pediatrician or family physician or community health department, and that would be that. That information would then go on their health records, and they'd go on with their more moral lives. Let us also suppose that if the program were administered covertly it would be conducted in similar fashion. When children are scheduled to receive vaccinations, they are at the same time given the moral bioenhancement, but neither the children nor their parents or guardians are told about the moral bioenhancement and it doesn't go in their health records. The administration of it could be double- or even triple-blinded, so that only a few individuals are aware of the moral bioenhancement. Everyone would go on with their lives unaware of the moral bioenhancement.

[13] Grill, K., & Dawson, A. (2015). Ethical frameworks in public health decision-making: defending a value-based and pluralist approach. *Health Care Analysis*.

The question is: which is the most ethically desirable scenario? I argue it is the second scenario in which the moral bioenhancement is administered covertly.

4.1 Values in public health ethics

The task is to compare the two methods of administering the compulsory moral bioenhancement and see which one is the most ethically permissible. Consider first the fact that as compared to a covert moral bioenhancement program that is blind to everyone except few, an overt program would reduce the expected utility of the program. It would reduce the expected utility of the moral bioenhancement program because if people knew that they were being morally bioenhanced, at least some of them would fail to receive the bioenhancement. They would request exemptions from the policy on the grounds that it conflicts with their religion or their personal convictions, or they would falsely believe that the moral bioenhancement leads to various disorders or diseases unrelated to the intervention. People would slip through. Some would slip through because of failing to pay attention, while others would outright refuse the intervention. That this would happen is obvious when we consider policies on vaccination or quarantine: people refuse vaccines or otherwise fail to get them, and people slip through quarantines and other methods of isolation.

If the moral bioenhancement were overt, the expected utility would be less than it would be if it were covert. It's not that the utility of preventing ultimate harm is less; it's that the expectation that the moral bioenhancement will succeed in preventing it is lower. The more people that slip through the compulsory moral bioenhancement, the lower the expectation that ultimate harm will be prevented. If the program were covert, people would be unaware of the intervention, and so would not be in a position to avoid it, resulting in many fewer people failing to receive the intervention.

Both overt as well as covert compulsory moral bioenhancement programs would restrict the range of moral attitudes, dispositions, and behaviors of its participants. The range of moral attitudes, dispositions, and behaviors that would be restricted would be the same for both types of program, as it is the intervention upon these that is presumably

necessary to prevent ultimate harm. So the extent to which the interventions themselves are liberty-restricting, the liberty-restrictions will be equal between a covert and an overt program. But for overt compulsory moral bioenhancement programs, participants would also *know* that their moral attitudes, dispositions, and behaviors are being intervened upon. For some of these people who know that their moral capacities are being restricted, they will desire to not be so restricted. Thus, the desires of these people will be frustrated, which results in suffering.

If the program were covert, the people who desire to not have their moral capacities restricted wouldn't be aware of any restriction, so from their perspective the desire to not be restricted wouldn't be frustrated, which means they wouldn't suffer from knowing that they are participating in a compulsory moral bioenhancement program.

This point rests on the idea that, all things considered, there is disutility in not just having desires frustrated, but also in *knowing* or *believing* that they are frustrated.[14] The same point could also apply to other public health programs, such as those that require people be vaccinated. Some people desire to not be vaccinated. When these people knowingly receive a vaccination—to attend school, for example—their desires are frustrated, and this frustration causes suffering. If it were possible to achieve all of the benefits of vaccination without having to cause the suffering that results from believing that one is vaccinated, then that would be preferable to actual vaccination procedures. To put it another way, in comparing an overt program with a covert program, the disutility of an overt program is greater, because in addition to the disutility of the fact of desires being frustrated, there's also the further disutility of believing that they are frustrated. This latter disutility is absent in the case of a covert program.

The frameworks of public health ethics require balancing the expected utility with the promotion or demotion of other values. One of these values is liberty. A covert compulsory moral bioenhancement program is less liberty-restricting than a similar overt program is. The discussion above points out that if the moral bioenhancement program

[14] This is not to say that having unconscious desires be unconsciously frustrated doesn't matter. I am merely pointing out that in an overt program, there is an additional awareness of desire frustration.

were overt, inevitably some people would refuse or otherwise fail to receive the intervention. Because the program is compulsory, however, policies would be required to compel such people to undergo the intervention. These policies would take the form of isolation (e.g., preventing dissenters from fully participating in society), taxes or fees as penalties, or, in severe cases, imprisonment. All of these methods of compulsion restrict liberties.

Moreover, given that the expectation of preventing ultimate harm is lower for an overt program, the potential for more significant liberty restrictions is greater, as our liberties may be more likely to be restricted by our harsher environments that result from having undergone ultimate harm. And upon one's death from ultimate harm, one's liberties are fully restricted—dead people have no liberties. These liberty restrictions should not be overlooked. The world in which ultimate harm has occurred is a world in which people have very few liberties. Compared to a covert program, an overt program makes this world more likely.

If the program were covert, however, enforcement of liberty restrictions would be unnecessary, as people would be unaware of the intervention in the first place and so there would be no need for such policies to compel participation. And the potential for the liberty restrictions that ensue from having undergone ultimate harm is also lower.

Some may think that utility and liberty are the most important values to consider in implementing a public health program. Balancing the liberty and utility of a covert program and an overt program, the better balance is that of a covert program, because it has greater expected utility and is less restrictive of participants' liberties.

But even when other values are balanced with the utility of preventing ultimate harm, a covert program is preferable. A covert program better promotes equality, because by keeping the program covert to everyone, the program ensures that all participants are treated equally. It is totally impartial. In an overt program, it would remain open that some populations are in a better position to avoid the intervention, such as those that could easily afford the penalties imposed for refusing, or those that do not rely on public health clinics.

Another potential source of unequal treatment is that likely many physicians would disagree with the policy, putting them in a better position to refuse to administer the moral bioenhancement. Based on this variance of attitudes within physicians, it is likely that the treatments would be administered unequally.[15]

Similarly, a covert program would be fairer than an overt program. Because everyone would receive the moral bioenhancement, there is no population that would be forced to bear a disproportionate burden. Some populations may bear a greater burden, such as psychopaths whose moral psychology must change to a greater degree than others.' For these populations, the burden of the moral bioenhancement would be greater (similar to a higher tax), but it is not a disproportionate burden. An overt program, however, may encourage others to find ways to avoid receiving the enhancement, meaning that they wouldn't be required to bear any burden, which is unfair.[16]

A covert program would also better promote population health than an overt program. As discussed above, the expectation that ultimate harm will be prevented is lower in an overt program, so the threat to population health is greater. For the same reason, an overt program may not promote individual health as much as a covert program would. Otherwise, the two programs promote population and individual health equally.

[15] A potential source of inequality in a compulsory program is between the people administering the program and the people receiving it. But the potential for this inequality is due to the program being compulsory, not to it being covert. Overt programs have the same potential to introduce this kind of inequality.

[16] Some authors have complained that moral bioenhancement would be inegalitarian (e.g., Sparrow, R. Egalitarianism and moral bioenhancement. *The American Journal of Bioethics : AJOB*, *14*(4), (2014) 20–8.). In addition to the replies to this claim (Lechner, S. Why moral bioenhancement is a bad idea and why egalitarianism would make it worse. *The American Journal of Bioethics : AJOB*, *14*(4), (2014) 31–2; Persson, I., & Savulescu, J. Against fetishism about egalitarianism and in defense of cautious moral bioenhancement. *The American Journal of Bioethics : AJOB*, *14*(4), (2014) 39–42; Wilson, A. T. (2014). Egalitarianism and successful moral bioenhancement. *The American Journal of Bioethics : AJOB*, *14*(4), (2014) 35–6), a covert program would not be inegalitarian.

The promotion of other values may also turn out to be equal between overt and covert programs. Solidarity would presumably be the same between the two types of programs, though it is possible that it would be greater in an overt program: the administration of the program could induce people to join each other in protesting it. And, if revealed, a covert program would undermine social trust.

Indeed, the disclosure of a covert program may undermine social trust so much that it could destabilize the society and its government. In this respect, the potential for such destabilization is much lower in an overt program. However, in the event of ultimate harm, society is also destabilized. I claim above that an overt program is less likely to prevent ultimate harm. So, although in one respect an overt program makes destabilization less likely, in another respect it makes it more likely. It's therefore not clear that an overt promotes social stability more than a covert program does, though it may better promote social trust.

Another value is that of transparency. Obviously, an overt program will better promote this value.

Comparing the two programs, a covert program wins on the values of utility, liberty, equality, and fairness. They are approximately equal in promoting health. An overt program wins on the values of promoting transparency and solidarity, and potentially trust.

What about the values of honoring duties and respecting autonomy? I assume that many readers are uncomfortable with the notion of a compulsory covert moral bioenhancement program, and that the source of this discomfort is the judgment that such a program undermines a person's autonomy, which we have a duty to promote. This, I assume, constitutes the primary reason that one would object to the proposal of making the compulsory moral bioenhancement covert. I address this objection shortly.

4.2 Conditions in public health ethics.

In addition to the values discussed above, there are plausible conditions that need to be satisfied in order for a liberty-restricting public health program to be justifiably implemented.

One of the two foundational assumptions is that a moral bioenhancement program is necessary. But the assumption does not distinguish between a covert program and an overt program. Thus, it is best to say that both types of programs would satisfy the condition that the program be necessary. Another assumption is that the compulsory moral bioenhancement is effective and, again, the assumption makes no distinction between the effectiveness of a covert program versus the effectiveness of an overt program. So both satisfy the condition that the intervention be effective.

A third condition is that the program be proportional to the threat. Whether a covert or overt program is more proportional to the threat depends on the burdens the programs require subjects bear. As I argue above and below, overt programs require that those receiving the enhancement bear greater burdens. Thus, if I am right about these other values and conditions, then a covert program is more proportional to the threat of ultimate harm than is a covert program, because it burdens subjects less.

The covert and overt programs diverge with the next two conditions: that the program be the least restrictive means and that it be transparent. As argued above, a covert program restricts liberties less than an overt program does. So an overt program doesn't satisfy that condition. But it does satisfy the condition that the program be transparent, whereas the covert program doesn't.

Thus, as measured by the collection of values and conditions that public health ethics frameworks propose, the two types of programs are similar. But the differences are key. A covert program better promotes or preserves utility, liberty, equality, and fairness. It fails to be transparent, promote trust, or promote solidarity.

Whether one thinks that a covert program is preferable to an overt program depends on whether one thinks that the values of utility, liberty, equality, and fairness are more valuable than transparency and the promotion of trust and solidarity. I believe they are, and for this reason a covert program is preferable to an overt program. I offer no criteria for ranking these values. But utility, liberty, equality, and fairness have enjoyed a

more central role in moral and political philosophy than have the values of transparency, trust, and solidarity. So those wishing to claim that an overt program is ethically preferable have the burden of establishing that transparency, trust and solidarity are more valuable than utility, liberty, equality, and fairness.

4.3 Implementation

I have argued that a covert moral enhancement program is preferable to an overt program, and have done so on ethical grounds. Though the argument is valuable even if it is impossible for a covert program to be implemented, the practical implications are much greater if such a program can be implemented.

It is admittedly difficult to provide examples of other public health interventions of the sort a covert moral enhancement program would be. If I know about them, then they weren't very covert, and so wouldn't be good examples of how a covert moral enhancement program should be administered.

But there may be ways to administer the moral enhancement to the relevant populations, without those populations knowing about the enhancement. The method of delivery would depend on the mechanism of action of the enhancement, but one possible way of distributing it to the relevant populations is by way of the public water supply. Another way it could be distributed is by packaging it with various vaccines, while eliminating most exceptions. Or perhaps it could be distributed through forced air systems in public buildings, or some combination of these.

There are of course members of the population that don't use public water, don't get vaccinated, and don't visit public buildings. For example, there are many people in rural areas who get all of their water from wells drilled into natural aquifers. These people wouldn't regularly come into contact with the public water supply. But they would still spend time in public places or get vaccinated, both opportunities to receive the enhancement.

There are still individuals who never use public water, never get vaccinated, and never go out in public. And if the enhancement were only administered through these

channels, such individuals would never receive the enhancement. The threat of ultimate harm is greater now than it has been because people now have greater access to information and materials that are capable of triggering ultimate harm, such as advancing technologies or widespread use of fossil fuels. People who refrain from public services may not have widespread access to the information and materials that make ultimate harm a threat. There are large segments of the human population who lack access to utilities, vaccinations, or public services of any kind. And if these are the channels of distribution, then these segments of the population will not receive the enhancement.

But these segments of the population aren't the ones who increase the threat of ultimate harm—they not only lack access to public services, but also to the information and materials that make ultimate harm a threat.[17] In some cases individuals in these segments could travel to places where the information and materials are accessible, but in so doing they would be traveling to a place with, and very likely use, public services, and would receive the enhancement. It's the bad actor with access to a published recipe for bird flu and the intelligence and infrastructure to produce it or the person who consumes the resources which contribute to climate change that we have to worry about, rather than a member of a recently discovered Amazonian tribe or someone who never engages with the public.

Though the methods of distribution described above occur by way of public utilities or public spaces, there is no implication that the administration of a covert moral enhancement must be a matter of public policy. Governments and bureaucrats are possible administrators, but the argument doesn't require that a covert program be carried out by any particular individual or group. A covert program would indeed require excellent coordination and control to not only maximize the number of people who receive enhanced moral capacities but also to maintain secrecy. Such coordination is not logically impossible, not metaphysically impossible, and not even practically impossible. It may even be easier to administer than other attempts at secrecy, if the administrators

[17] For example, access to "dual-use" research, which is research that can be used for public good or public harm. In some cases, access only requires internet access.

themselves are being enhanced, as their moral motivation to maintain secrecy may be stronger than it otherwise would be. Keeping a covert program covert would be a challenging obstacle. But just because it would be a challenging obstacle doesn't mean that my argument is unsound.

Further, if one is concerned about the difficulty of implementing such a program, and for this reason doubts my argument, consider that it is also no easy task to cobble together a range of public programs which individually are aimed at preventing one potential avenue of ultimate harm but collectively aimed at preventing ultimate harm. It's not clear that would be any easier to implement. Such a collection would need to include programs aimed at stopping or reversing climate change, programs preventing the proliferation of nuclear arms and other devastating weapons along with the information used to manufacture them, programs protecting democratic institutions from electing powerful bad actors, programs preventing dictators from becoming powerful bad actors, programs preventing scientists from developing, even accidentally, threatening materials or organisms, programs that prevent amok AI, programs that allow humans to colonize space before the death of the sun, etc.

It seems reasonable to me that a covert program can be implemented. Others may not be convinced. But this disagreement isn't about whether it's true that if moral bioenhancement ought to be compulsory, it ought to be covert. One could instead argue on the basis of a covert program's impossibility that it's false that a covert, and therefore compulsory, moral bioenhancement ought to be carried out. This argument requires that it be true that ought implies can. The ought implies can principle is that one ought to do something, only if one is able to do that thing. It's controversial that the principle is true. But if it is true, and if a covert program is impossible, then it's false that moral bioenhancement ought to be covert (by the ought implies can principle). And if it's false that moral bioenhancement ought to be covert, then it's false that it ought to be compulsory (by the truth of the conditional I argue for). So, the impossibility of covert moral bioenhancement doesn't refute my argument that if moral bioenhancement ought to be compulsory, then it ought to be covert (because the falsity of a consequent doesn't

entail the falsity of the conditional). But it could refute that moral bioenhancement ought to be compulsory, by way of the principle that ought implies can.

5. Objections

5.1 It's unrealistic

Above I warn that some readers may find the assumptions to be so demanding that they push the argument into the unreal. What is far more likely is that the intervention is not fully safe or fully effective. What happens when we can only be eighty percent sure that it will be effective, or when ten percent of those receiving the enhancement will suffer a terrible side effect?

Though the certainty of whether the intervention will be safe or effective impacts how the different values are balanced, I don't think the impact is significant enough to drastically alter the balance. Suppose there's an equally low chance that an overt and covert program are effective in preventing ultimate harm, say ten percent. In that case, as long as it is necessary to prevent ultimate harm, then a covert program is preferable. If there's a one percent chance, or even substantially lower, of it being effective, then ultimate harm may be overdetermined, in which case no program should be instituted. If an overt program were considerably more likely to be effective, or considerably safer, then it would be preferable. But if the safety and effectiveness of an overt program is similar to that of a covert program, and a moral bioenhancement program is necessary to prevent ultimate harm, then a covert program is preferable.

Additionally, some may find the possibility of ultimate harm remote, and if not remote, at least not a threat that will emerge in the near future. Ultimate harm is not inevitable—humans may die away slowly and quietly, without experiencing the significant suffering and loss the prevention of which justifies moral bioenhancement. But there are many potential avenues of ultimate harm: bioterror, nuclear war, devastating climate change, meteor strikes, and geologic disasters are just a few possibilities. At the very least, if the human species (or its evolutionary descendants)

survives to the death of the sun, ultimate harm is likely, unless we colonize space, prevent the death of the sun, or become disembodied.

Given the difficulty in determining how likely ultimate harm is, it's difficult to say whether a covert moral bioenhancement program will ever need to be administered. What likelihood of ultimate harm is sufficient to trigger the program? A fifty percent chance of ultimate harm seems sufficient to trigger it, but a one tenth of one percent chance seems insufficient, though humans often go to great lengths to prevent unlikely events, even when the costs associated with those events are comparatively low.

5.2 It's dishonest

If preventing ultimate harm requires lying, even by omission, then lie we must. Even if the dishonesty is not necessary *and* it is unambiguously wrong to be so dishonest, the wrongness of the dishonesty does not outweigh the utility of preventing ultimate harm.

To argue that it is not necessary to be dishonest about the moral bioenhancement is nothing more than the claim that a covert moral bioenhancement is not necessary for the prevention of ultimate harm. However, I am not claiming that a covert moral bioenhancement program is necessary. I am arguing that it is preferable to an overt program.[18] To measure this preference, I have appealed to frameworks of public health ethics. And it seems according to these frameworks that a covert program is preferable.

One may also think that there is disutility in dishonesty. Thus, when do our balancing of values, we must include this disutility that comes along with a covert moral bioenhancement program. The alleged disutility of the dishonesty would have to outweigh the risk of ultimate harm. It seems obvious to me that it doesn't, that it is more akin to a little white lie when compared to the potential risk of being honest. But others may think that it's worth the greater risk to the existence of humanity to be honest.

[18] If moral bioenhancement is necessary, the program can be either overt or covert.

5.3 It's unlike other public health programs

Compulsory moral bioenhancement is a matter of public health. According to this objection, other public health programs, in particular those that require participation, are more precisely targeted at those who need them. For example, only children are required to get vaccinated, because they are the people who need the intervention. Since not everyone needs moral bioenhancement, its implementation is disanalogous to other public health programs.

While it's true a covert program would be unlike anything we know about (though it could be like programs we don't know about), it's not true that other compulsory public health programs only target those who need the intervention. Not all children need vaccinations for their own health—some might never get the flu or chicken pox or polio in the absence of the vaccine. But making the vaccine compulsory, even in the absence of individual need, protects the public. Quarantines routinely isolate people who have no indication of disease. People who brush their teeth regularly and abstain from sugary foods don't need fluoridated water. Some people may not need moral bioenhancement, (though I suspect this number is very low, as most of us contribute to various potential catalysts of ultimate harm). An individual's lack of need of a public health intervention does not imply that requiring that individual's participation in the intervention is unjustified.

5.4 It's paternalistic

The prevention of ultimate harm is of far greater value than the prevention of paternalistic policies. And in any case, whether the compulsory program is covert or overt is independent of whether it is also paternalistic. In other words, an overt program is equally paternalistic.

5.5 It diminishes autonomy

Some may think that by intervening upon a person's moral attitudes, dispositions, and behaviors the moral bioenhancement program is preventing a person from expressing their autonomy. It prevents him or her from acting contrary to the intervention. And if that's what the person wants to do, then the intervention prevents the person from doing what he or she wants to do, when, in the absence of the intervention he or she could. This diminishes the person's autonomy. And since we have a duty to respect autonomy, a covert compulsory moral bioenhancement is not ethically permissible. However, if the compulsory moral bioenhancement program were overt, at least then the person would be in a position to accept or reject the intervention. Or so the objection goes.

The notion of autonomy that I have in mind says to be autonomous is to be in a position to self-govern in accordance with one's desires, values, or other attitudes. This may not capture every conception of autonomy, but it at least seems central to the different variations.[19]

With this conception of autonomy in hand, it may be that a covert moral bioenhancement program does not diminish a person's autonomy at all. But if it does, it at least diminishes it less so than an overt program. If to be autonomous is to be in a position to self-govern in accordance with one's desires, values, and other attitudes, then one's autonomy could be diminished, if one is out of position to self-govern. Thus, to be autonomous one must have the appropriate power to reason and choose, but must also be able to reason and choose independent of external manipulation, and one's desires, values, and attitudes must be authentic.[20] If this is right, then the objection that a covert moral bioenhancement program violates a person's autonomy amounts to the claim that it violates either one's independence from external influence or the authenticity of one's desires, values, and other attitudes.

---

[19] Christman, John, "Autonomy in Moral and Political Philosophy", The Stanford Encyclopedia of Philosophy (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2015/entries/autonomy-moral/>.

[20] Christman op. cit. note 18

Consider first the independence condition, the condition that one's reasoning and decisions be free of external influence. It is common that our reasoning and decisions are dependent upon external influence. For example, experimental evidence suggests that a person's moral judgment about which life-saving public health program to implement in response to a disease outbreak is dependent on the words used to describe the programs.[21] This sort of framing effect on our moral judgments is common.[22] A slightly different type of influence is that of one's environment. In one experiment, subjects finding a dime in a phone booth were much more likely to mail a seemingly misplaced addressed envelope than those who found no dime.[23] And in another study subjects were more likely to help an injured man if the ambient noise was normal than they were if a loud motor was running.[24]

In these experiments, and presumably all of the instances to which they generalize, the subjects arrive at a moral judgment induced by others without awareness of the influence. Intuitively, these types of external influence do not violate the independence condition. If they do violate the condition, then many, if not most, of our moral judgments are not autonomous, and moral bioenhancement doesn't constitute a unique threat to autonomy. As long as the moral bioenhancement exerts a similar type of

[21] Tversky, A., & Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science, New Series*, *211*(4481), 453–458.
[22] Sinnott-Armstrong, W. (2008). Framing moral intuitions. *Moral Psychology Vol. 2.* MIT Press.

[23] Isen Franklin, A. M., & College Paula levin, m. F. (1972). Effect of feeling good on helping. *Journal of Personality and Social Psychology*, *21*(3), 384–388

[24] Mathews, K. E., & Canon, L. K. (1975). Environmental noise level as a determinant of helping behavior. *Journal of Personality and Social Psychology*, *32*(4), 571–577.

influence, the fact that a moral bioenhancement program exerts external influence makes it no different from these other ways our judgments are influenced.[25]

Even if a compulsory moral bioenhancement program does undermine our independence from external influence in an autonomy-violating way, it's the compulsory nature of the intervention that does it, because it's the intervention that is the external influence. This intervention, in a compulsory moral bioenhancement program, occurs whether the program is overt or covert. So, if a covert moral bioenhancement program undermines a person's independence from external influence, then so does an overt program. And if they violate it (or don't) equally, then autonomy violation is not a reason to prefer one program over another.

However, an overt program and covert program are not equal when it comes to the potential violation of the authenticity condition, or the condition that one embrace one's desires, values, and other attitudes as one's own. In this case, a compulsory moral bioenhancement program *does* violate autonomy, but only if the program is *overt*. If a person is compelled to participate in a moral bioenhancement program, and the person believes that the new moral capacities—including the new desires, values, and other attitudes—are caused by the enhancement, it is much more difficult to see how the person would embrace these capacities as their own.[26] The knowledge that some of one's moral capacities are the result of manipulation by another agent undermines trust in their authenticity. Thus, an overt program is likely to violate the authenticity condition.

If the moral bioenhancement is covert, the person is in a much better position to embrace the new capacities as one's own. Though the new capacities are in fact not one's own, there are fewer obstacles to embracing them as one's own, such as the knowledge that they are not. If from the person's perspective he or she cannot notice any difference

[25] Whether it exerts a similar influence would depend on the mechanism of action. One may think that even when one's moral judgment is influenced by the situation one finds oneself in, one still has the ability to reject the influenced moral judgment, but that in the case of moral bioenhancement one would not be in such a position. This misses the lesson of framing effects—they take us out of the position to accept or reject the framed moral judgments.

[26] Author discusses the implications of knowing one's been morally enhanced in detail in (Redacted).

between one's own capacities, and the new, enhanced capacities, then from their perspective there would be no reason to not embrace the new capacities as one's own.[27] So, as long as the enhancement is covert and the person doesn't believe he or she has been enhanced, the program introduces no additional threat to the embrace of one's capacities as one's own. But an overt program does. So, if a moral bioenhancement is compulsory, to best preserve authenticity, it is preferable for the program to be covert.

Even if a moral bioenhancement program does diminish a person's autonomy, there is no implication that to do so is wrong. Quarantine diminishes the isolated person's autonomy, and we don't think that quarantine is wrong. The same goes for forced treatment of people with mental health disorders disposing them to be harmful to themselves or others. I am arguing that we should treat those that are disposed to engender ultimate harm as hazards to the public's health. So, if moral bioenhancement of this population is going to be compulsory like other public health interventions, it ought to be done covertly, as this best preserves or promotes the values and conditions of frameworks of public health ethics.

6. Conclusion: Should Moral Bioenhancement be Voluntary?

I have argued for the truth of the conditional that if moral bioenhancement ought to be compulsory, then it ought to be covert. This conditional is a potential premise in a variety of different arguments. It could be conjoined with an affirmation of the antecedent (which I have not attempted here), arriving at the conclusion that moral bioenhancement ought to be covert. Alternatively, one might think that if compulsory moral bioenhancement implies covert moral bioenhancement, then so much the worse for compulsory moral bioenhancement. That is, the conditional I argue for could be conjoined with a denial of the consequent (which would also need to be independently supported, for example by arguing that covert moral bioenhancement is impossible *and*

[27] Author (redacted) addresses what the enhancement needs to manipulate in order for a subject to fail to notice it.

that ought implies can), arriving at the conclusion that moral bioenhancement ought not be compulsory.

From here it is short step to voluntary moral bioenhancement. Either moral bioenhancement should be voluntary or compulsory; since it shouldn't be compulsory, it should be voluntary. Voluntary moral bioenhancement has been supported by several authors. Vojin Rakić argues that voluntary moral bioenhancement is preferable to compulsory moral bioenhancement on the grounds that compulsory moral bioenhancement undermines humans' free will.[28] Even if we grant that the prevention of ultimate harm motivates moral bioenhancement (there may be other grounds, according to Rakić) the undermining of free will is a type of ultimate harm in and of itself.[29] Furthermore, Rakić argues that making some types of moral bioenhancement (e.g., administration of oxytocin) compulsory renders the enhancement process ineffective, because making it compulsory undermines moral reflection. And the potential benefits (those that would outweigh the costs associated with making it compulsory) of oxytocin can be achieved only when accompanied by moral reflection.[30] Similarly, John Harris argues that moral bioenhancement may undermine our freedom, and in so doing undermine our virtue.[31]

---

[28] Rakić, op cit note 8.; Rakić, V. (2014). Voluntary moral enhancement and the survival at-any-cost bias. *Journal of Medical Ethics*, *40*(4), 246–250; Rakić, V. (2017). The issues of freedom and happiness in moral bioenhancement: continuing the debate with a reply to Harris Wiseman. *Journal of Bioethical Inquiry*, *14*(4), 469–474.

[29] Intereestingly, Rakić (2017, p. 70), op. cit. note 27 offers a parenthetical comment that what voluntary moral bioenhancement really preserves is the *perception* of freedom. A covert compulsory program would also preserve that perception, as I argue for here and in my (redacted)

[30] Rakić, V. (2017). Compulsory administration of oxytocin does not result in genuine moral enhancement. *Medicine, Health Care and Philosophy*, *20*(3), 291–297.

[31] Op. cit. note 4

As I argue above, a compulsory *covert* moral bioenhancement does not undermine our autonomy. Though Harris Wiseman casts doubt on the usefulness of the distinction between voluntary and compulsory moral bioenhancement,[32] if there is such a distinction I have no issue with the argument that concludes that moral bioenhancement ought to be voluntary. My argument cuts both ways, depending on what additional premise one supports. One person's *modus ponens* is another person's *modus tollens*. I have addressed at length elsewhere (Author, redacted) the idea that knowing that one has been morally bioenhanced, knowledge which volunteers would possess, epistemically defeats the enhancement, rendering it ineffective. Further, if moral bioenhancement is necessary to prevent ultimate harm, I'm skeptical that voluntary moral bioenhancement is up to the task. People who volunteer for moral bioenhancement are not those about whom we should be most worried—it's those who have no interest in being better people that should worry us.[33]

This is not to say that the only alternative is to pursue voluntary moral bioenhancement. We may, as Harris suggests, opt to forego moral bioenhancement altogether, favoring instead cognitive enhancement, which he claims is sufficient for moral enhancement.[34] One could argue from the conclusion that compulsory moral bioenhancement implies covert moral bioenhancement to, a few steps later, the rejection of moral bioenhancement altogether, if one can offer independent reason to reject voluntary moral bioenhancement. [35] Depending on the other premises, it could be that since compulsory moral bioenhancement implies covert moral bioenhancement, the best alternative is no moral bioenhancement at all. Thus, the implication for which I argue can

---

[32] Wiseman, H. (2014b). SSRIs and moral enhancement: looking deeper. *AJOB Neuroscience*, *5*(4), W1–W7

[33] Wiseman makes a similar point, op. cit. note 29.

[34] Op. cit. note 4. Rakić (op. cit. note 28) argues that relying on cognitive enhancement is not sufficient for moral enhancement, as cognitive enhancement does not necessarily lead to morally enhanced behavior.

[35] Thanks to an anonymous referee for bringing this argument to my attention.

be used in a number of arguments, potentially cutting the ethics of moral bioenhancement at its joints.