

An Evaluative Conservative Case for Biomedical Enhancement

By John Danaher

Forthcoming in the *Journal of Medical Ethics*

It is widely believed that a conservative moral outlook is opposed to biomedical forms of human enhancement. In this paper, I argue that this widespread belief is incorrect. Using Cohen's evaluative conservatism as my starting point, I argue that there are strong conservative reasons to prioritise the development of biomedical enhancements. In particular, I suggest that biomedical enhancement may be essential if we are to maintain our current evaluative equilibrium (i.e. the set of values that undergird and permeate our current political, economic, and personal lives) against the threats to that equilibrium posed by external, non-biomedical forms of enhancement. I defend this view against modest conservatives who insist that biomedical enhancements pose a greater risk to our current evaluative equilibrium, and against those who see no principled distinction between the forms of human enhancement.

Keywords: Evaluative Conservatism; Biomedical Enhancement; Automation; Technological Progress

1. Introduction

Conservatives like things the way they are; enhancement technologies allow us to change them. It is no surprise then that there is a widespread belief that conservative moral principles lead to the rejection of human enhancement. In this article, I argue that this widespread belief is wrong in a number of critical respects. I argue that the promotion of biomedical human enhancement could be crucial if we are to conserve our current evaluative equilibrium against an array of technological threats. I defend this argument in four parts. First, I situate my argument within both the conservative moral tradition and the enhancement debate. Second, I present and clarify the argument itself. Third, I defend the central premise of the argument, noting three specific ways in which biomedical enhancement could help to preserve our current evaluative equilibrium. And finally, I respond to six objections to this argument.

2. Enhancement and Evaluative Conservatism

It is important to clarify the style of conservatism adopted throughout the remainder of the article and to explain how the argument I defend it is situated relative to other, somewhat similar, arguments.

Conservatism is a rich philosophical tradition. My argument draws on a thin, but core, version of this tradition. The version has two components: an evaluative one and an epistemic one. The evaluative component maintains that we should seek to hold onto or preserve existing sources of value *even if we are promised newer better sources of value in their stead*. This evaluative conservatism is taken from the work of Jerry Cohen.[1, 2] Cohen's conservatism holds that there is reason to favour an existing source of value (V_1) over a potential/future value (V_2), even if there is also reason to suspect that the (expected) value of V_2 would be greater than V_1 . An example helps.[3] Suppose you currently have a pet dog – call him Fido. You have established a deep and affectionate relationship with Fido. One day someone offers you a new dog – call him Rex. You are assured (and you believe) that you could have an even deeper and more affectionate relationship with Rex, but only if you give up Fido. Presented with such a choice, many people would hang onto Fido. Why? Because they are evaluative conservatives: They wish to continue to value that which is valuable and valued by

them,¹ partly because the thing is itself a source of value and partly because of the personal and/or cultural history they associate with that source of value. In other words, they want to hold onto existing sources of value because of how those sources figure in their personal and cultural identity, in preference to alternative and arguably better sources of value. My conservative case for biomedical enhancement works from this evaluative form of conservatism. It maintains that we should favour biomedical enhancement because doing so will help us to preserve existing personal, social and political sources of values that are currently under threat. In doing so, it also appeals to our personal and cultural attachment to these sources of value.²

In addition to the evaluative component, I include an epistemic component in the conservative stance adopted in this article. This helps to situate my argument within the popular Burkean tradition of conservatism.[4] This form of conservatism is sceptical of change on the grounds that existing social arrangements can embody a degree of distributed knowledge that is difficult to make accessible to conscious awareness and deliberation, and hence there is significant evaluative risk in seeking to change from these existing arrangements. In this sense, the argument presented could be linked with the concerns expressed in debates about the role of moral risk in practical reason, though I will not spell out those links in what follows.[5-8] The epistemic component is relevant because although Cohen's evaluative conservatism works even in a situation in which we know for sure that the future source of value will be better than what we currently have, it seems more realistic to assume that we make our moral choices in the face of risk and uncertainty.

I make no further conservative commitments in this article. I would suggest, though I will not defend this suggestion, that the evaluative and epistemic components that I have identified represent the core of most forms of political and moral conservatism. But to this core additional commitments are often added. For instance, bioconservatives often make strong assumptions about the fixity and imperfections of human nature/biology and the desirability of a natural order.[9] I deliberately eschew

¹ Pugh et al [3] discuss the distinction between valuing the valuable and valuing the valued. The former is about intrinsic value; the latter is about personal history with the source of value.

² A reviewer pointed out that evaluative conservatism could have paradoxical consequences. What if, for example, one currently values constant change? If so, then one would constantly favour changes to the existing status quo, which hardly seems conservative. Interestingly this is a criticism of enhancement that some modest conservatives develop (e.g. Hauskeller 2011). This might be a logically permissible but it seems unlikely that anyone espousing a conservative moral outlook would favour this view hence I ignore it in what follows.

those commitments here. I do so for two reasons. The first is that these thicker and more traditional forms of conservatism have been widely defended and critiqued already in the debate about human enhancement.[9-11] The thinner and more modest forms of conservatism are more novel and underexplored. Second, the thinner form of conservatism is arguably more formidable: by being modest in its conservative commitments it tries to show how conservative values can appeal across the political and moral spectrum. Cohen, for instance, was no political conservative but he found evaluative conservatism appealing.

Proponents of this thinner form of evaluative conservatism have used it argue against human enhancement technologies. Cohen believed, and others have agreed, that evaluative conservatism provided powerful reasons to oppose enhancement.[3] Likewise, recent anti-enhancement authors such as Nicholas Agar and Michael Hauskeller,[12-15] although not explicitly calling themselves conservatives, have challenged the enhancement project on the grounds that it threatens existing sources of value. Agar worries that enhancement technologies prioritise hypothetical external goods over existing internal (and species-specific) goods. Hauskeller explicitly laments the fact that enhancement favours the possible better over the existing good.[14] My goal in this article is to suggest that evaluative conservatism actually provides powerful reasons to favour human enhancement. It does so because enhancement can help us to preserve much of what is currently valuable about our way of life.

This basic insight is not novel. Pugh, Kahane, and Savulescu[3] have pointed out that Cohen's conservatism can be used to defend some forms of enhancement. For instance, they argue that life extension therapy can be defended on the grounds that it preserves the value that attaches to our current lives and, similarly, mood enhancement drugs can be defended on the grounds that they allow us to preserve the value of existing interpersonal relationships. Buchanan makes similar claims, pointing out that '[t]o avoid losing some of the good things we now enjoy, we will have to enhance ourselves in particular ways' ([11], p. 163). But none of these authors develops the conservative case beyond a mere suggestion. They shine a light down a possible pathway; they do not lead us to the end. That's what I want to do in what follows. I want to argue that evaluative conservatism can lead to the view that biomedical enhancement is close to an axiological imperative: something we must pursue lest we

lose much of what is currently valuable (and valued by us) in our political, economic and personal lives.

3. The Evaluative Conservative Argument

So the argument is that *biomedical enhancement* may be essential if we are to conserve our current *evaluative equilibrium*. This argument appeals to a number of controversial concepts, each of which needs to be clarified.

I start with the most obvious point of contention: the idea of *enhancement* itself. There are two prominent approaches to the definition of enhancement in the literature.[16 – 20] There is the *functionalist* approach, which holds that an intervention counts as an enhancement if it improves (or adds to) the functionality of a human being relative to some population-level (or species-level) norm.[9] And there is the *welfarist* approach which holds that an intervention counts as an enhancement if it increases the chances of someone living a good life, relative to their current circumstances.[16] The argument I defend works best with the functionalist definition, though it will also work with the welfarist definition, provided that the ‘good life’ is understood to include more than simply subjective well-being and, in particular, to include some sense of personal achievement, mastery and understanding.

The second point of contention is the focus on *biomedical* forms of enhancement. As Buchanan[9] notes, enhancement is a broad concept and could refer to interventions as diverse as farming, democratic governance, literacy, education, pharmacological drugs, brain stimulation, and neuroprosthetic devices. Some people insist that there are no principled distinctions between these forms of enhancement.[21] Others insist that there are.[22] In this article, I use the descriptor ‘biomedical’ to designate enhancement technologies that either directly target (e.g. pharmacological drugs) or directly integrate with (e.g. neuroprosthetics) our human biology. They are to be contrasted with external and automating technologies, which often make our lives easier, may indirectly increase our power to change the world around us, and so could also be said to have an enhancing effect. I will insist that there is an important distinction between the two types of technological intervention when it comes to the conservation of value.

This brings me to the third and final point of contention, which relates to my claim that biomedical enhancement can conserve our current *evaluative equilibrium*. I appeal to this concept in order to incorporate Cohen's evaluative conservative principles into my argumentative framework. I define 'evaluative equilibrium', loosely, to mean the set of values that undergirds and permeates our current way of life. In what follows, I particularly appeal to the set of values that undergird the dominant Western political and economic systems, and those that support our sense of personal achievement and satisfaction. Evaluative conservatism then enters the picture as an attempt to defend that current evaluative equilibrium in lieu of alternative evaluative equilibria promised or predicted by certain technological enthusiasts (even if those evaluative equilibria promise to be better). I appreciate that this notion of a readily identifiable evaluative equilibrium is questionable, but I hope to add some reasonably specific content to it below.

This brings us to the argument itself, which can be formulated like this:

- (1) We ought to maintain our current set of values (VS_1) over a potential/future set of values ($VS_2 \dots VS_n$), even if those future sets of values are potentially superior.
- (2) Prioritising biomedical enhancement can help us to maintain VS_1 and stave off threats to VS_1 posed by other technological developments.
- (3) Therefore, we ought to prioritise biomedical enhancement.

There is an unavoidable degree of empirical uncertainty in premise (2).³ This is why the argument talks about how biomedical enhancement 'can help' rather than 'will help', and why the conclusion is framed in terms of 'prioritisation' not actual use.

Nevertheless, the argument defends a reasonably strong view about the value of biomedical enhancement. It claims that the development of biomedical enhancement technologies could be a moral imperative due to the threat to our current evaluative equilibrium that is posed by other technological developments. Hence, it is not simply that biomedical enhancement is desirable or preferable. The strength of this should become clearer in the next section.

³ On the role of empirical uncertainty in the enhancement debate see Loi [23].

For the purposes of this argument, I assume that the conservative principle stated in premise (1) is correct. Readers will need to go elsewhere for its defence.[1-3] My focus thus switches to the defence of premise (2). I defend that over the remaining two sections of the article.

4. The Threat of Externalisation to our Evaluative Equilibrium

Proponents and opponents of enhancement often consider the ethical implications of enhancement technologies in isolation from other technological and social trends.⁴ Thus, for example, they might consider the impact of cognitive enhancement drugs among university students against the more or less static socio-normative background of contemporary (often Anglo-American) universities. This allows them to identify ways in which the use of such drugs may be desirable or undesirable, given that background. But it does not enable them to consider how the use of such drugs may fit into an evolving and ever-changing model of university life. The central theme of my argument is that this sort of static ethical analysis is a mistake. When considering the ethical implications of biomedical enhancement, we must take into account the dynamic and ever-changing socio-technical background in which such enhancements are likely to develop. It is when we do this that we see most clearly the conservative value of such technologies.

This means that the first thing we need to do in order to defend premise (2) is to consider the background of other technological changes against which the assessment of biomedical enhancement must take place. It is impossible to provide a complete description of that background, but we can at least highlight some of its most salient features. I would argue that the biggest and most important feature of that background is the rapid growth in computerisation, artificial intelligence, machine learning, robotics and other types of information technology. I don't view this as a particularly controversial claim. The growth in these technologies is readily transparent. It has been exponential for much of the past 60 years and some are convinced that this exponential growth will continue, and will have spillover effects as ICT 'infects' other technological developments.[24-26] There are three important manifestations of this trend (i) the rise

⁴ Some discussions do not omit this point. A good example is the work of Persson and Savulescu [63] on moral enhancement, which is explicitly framed in terms of broader climatological and technological changes.

of automating technologies in agricultural, manufacturing, professional, and service-related work;[27-30] (ii) the increasing use of data-mining and other automated decision-making technologies in governance systems;[31, 32] and (iii) the use of digital assistants and robot helpers in our personal lives.[32, 33] My claim is that these three technological developments pose a threat to our current evaluative equilibrium. Once we see this, we can see the conservative potential of biomedical enhancements. It is best to do this by addressing the threats individually.

I start with the economic threat. Most Western societies are characterised by mixed economic systems in which the government provides legal and administrative support for a system of private capitalistic enterprise. Most people are employed (or hope to be employed) within that system. These employments are a significant source of value to the individuals being employed and to society more generally. It is how goods and services get distributed to those who need or desire them. It is how individuals secure the economic rewards that make life possible and pleasurable. Many people tie their self-worth and self-actualisation to their success in their working lives, and are then socially evaluated on the basis of this success. Welfare systems provide some sort of safety-net, but this is usually insufficient and typically tied to the ability to work. Furthermore, welfare systems do not compensate for the non-monetary goods of work. As Ghaeus and Herzog have recently argued, work is a privileged context for the attainment of goods such as excellence/mastery, social contribution, sense of community and social status.[34] As such, working for economic reward forms a significant part of our current evaluative equilibrium. The problem with the rise of automation technologies is that they threaten to undermine this system of values. An increasing number authors are convinced that these technologies will displace much of the current workforce, thereby depriving them of this source of value.[27, 28] Furthermore, even those who are not convinced by the strong technological unemployment thesis seem to agree that automating technologies change the types of work that are available, making low-paid, precarious, but difficult-to-automate work more common.[35] Others even go so far as to draw the connection between this trend and the enhancement debate and argue that technological unemployment is likely to have a ‘disenhancing’ effect, at least in the welfarist sense of the term.[23]

Some welcome the disrupting effect of automation on work, but they are typically radicals who long for a ‘postwork’ future in which we have abandoned the ideological

glorification the work ethic.[36] There is little solace in this radical posture for the evaluative conservative who sees the current system as a source of value, one to which they have grown attached. The person who has worked all their lives in a profession like law and developed mastery over a certain set of valuable skills (e.g. legal research; document review; brief writing) is going to lose that source of value if all those skills can be automated. Even if the promised postwork future would be better for this person, the evaluative conservative would insist that the person is axiologically warranted in trying to hold onto what they current have. Others suggest that a basic income guarantee could address some of the problems posed by technological unemployment.[27, 28, 37] It would enable everyone to receive an income irrespective of their ability or willingness to work. This means some of the consumption-related value of work could be conserved. But the problem with this is that the non-monetary values that attach to work would not. So a basic income may partly address the concerns of the evaluative conservative, but not all of them. Biomedical enhancement, on other hand, is more hopeful. Biomedical enhancements could allow human beings to ‘keep pace’ with their automated rivals, either through the direct augmentation of biological systems, or through the grafting in of prosthetic devices. Biomedically enhanced humans could continue to be valued contributors to the economic system. This could help to conserve work as a source of value.⁵

Consider a simple example. A surgeon has dedicated their lives to mastering a set of highly precise motor skills and marrying that to a deep understanding of how the human body works. Their work is individually meaningful and socially valuable. Unfortunately, the hospital has decided to invest in robotic surgeons that are more precise, more efficient, less prone to fatigue and automatically updated with the latest medical knowledge. The surgeon is told that their work is no longer needed. The result might be improved medical outcomes for the hospital, but a significant loss in meaning and value for the surgeon whose identity is so closely tied to their competence at work. This outcome could have been avoided if the hospital (or whichever relevant authority) prioritized biomedical enhancement of human surgeons over investment in robotic surgeons. Human surgeons could be augmented with enhancing drugs that make their hands steadier, improve their concentration, reduce their fatigue and enable them to better absorb the latest medical knowledge. At least in principle, there was a way to

⁵ Loi [23] p 209 makes a similar observation in his discussion of technological unemployment, arguing that biomedical enhancement could used to counterbalance the disenchanting effect of workplace automation.

achieve better outcomes whilst still conserving the value of work to the surgeon. There were choices to be made about investment and prioritization of technological innovation that could have retained our existing work-related evaluative equilibrium.⁶

The second threat arises in the political sphere. Most Western societies are characterised by liberal democratic systems of governance. These systems of governance are undergirded by the commitment to moral equality and individual autonomy.[38] These commitments imply that governance structures are illegitimate unless they satisfy certain conditions.[38, 39] These legitimacy conditions come in two major flavours: instrumentalist and proceduralist. Instrumentalist legitimacy conditions justify governance structures in terms of their outputs, *i.e.* do they maximise some value beloved by the polity (less crime, more healthcare, more economic growth, longer lives, better education etc.)? Proceduralist legitimacy conditions justify governance structures in terms of the processes they use to produce their outputs, *i.e.* are these processes inclusive, participative, comprehensible, publicly accessible and so forth? Although some favour purely instrumentalist or proceduralist approaches to legitimation, most accept that a mix of instrumentalist and proceduralist conditions must be satisfied.

But here's the rub: algorithmic and automated decision-making is increasingly being used by political governance systems.[40-42] These systems limit or obviate the need for human input and oversight of decision-making. Algorithms can collect and organize decision-relevant information, draw inferences from that information, and use it to make important political, legal and bureaucratic decisions. Dormehl[32] gives the example of a facial recognition algorithm used by the Massachusetts Registry of Motor Vehicles to detect probable false identities in its database of driver ID photographs (157-8). Once the system found a suspiciously-similar match, an automatic letter was issued cancelling the driver's licence. Similar systems are being used to detect potential terrorist threats and likely tax cheats. And the number of such systems is rapidly growing.[43]

Their introduction is often defended on the grounds that they are more accurate, less costly, and more efficient than traditional bureaucratic systems.[32, 44] But are

⁶ There are also important political and social choices to be made regarding who benefits most from the productivity gains of either enhanced or automated labour. One major concern about automated labour is that it increases the share of wealth that goes to owners of capital and thereby exacerbates income and social inequality. While biomedical enhancement of workers is unlikely to solve inequality, it may be considerably *less worse* than automation since workers can at least retain their income from work.

they consequently more legitimate? This is highly debateable. They may be more instrumentally legitimate insofar as they can facilitate better outcomes, but their procedural virtues are dubious. These systems typically rely on mass surveillance, largescale data-mining operations, and machine-learning technologies. In addition to the privacy-related concerns that these systems raise, they also raise considerable transparency, comprehensibility and accessibility concerns. The highly technical nature of the systems makes them comprehensible to a relatively limited set of human elites. And more advanced versions of these systems are often impenetrable to even these elites, due to both the opacity of the technology, and the fact that they are woven into complex ecosystems of pre-existing algorithms.[45, 46] The result is that the increased use of algorithmic and automated systems of governance poses a threat to the procedural legitimacy of our political system, which again constitutes a central part of our evaluative equilibrium.

The threat can be explained by way of analogy. A long-standing concern in political theory is with the attractiveness of non-democratic governance. Some (e.g. Plato and JS Mill) have argued that human epistemic elites (i.e. humans with greater cognitive ability, education, technical knowledge and so forth) should be given more decision-making power. Others have challenged this on the grounds that deferring to human epistemic elites often means blocking ordinary people from participating in the decisions that affect their lives, and/or failing to explain and justify those decisions to them. The result is something Estlund calls the threat of ‘epistocracy’ – a threat that undermines participative democratic governance.[47] The growth of algorithmic governance gives rise to a similar threat: humans create algorithmic systems in the belief that they can gather and make sense of more information than humans ever could. They start deferring to the algorithms for epistemic reasons. This may be unproblematic at first, but as the algorithmic systems become more complex, comprehensive and opaque, the humans may end up deferring to them in blind faith. The result is something like a threat of ‘algocracy’ – humans are blocked from participating in and understanding the algorithmic decisions that affect their lives.[42]

The conservative impulse might be to address this problem by limiting, banning and regulating the use of such systems. This may have some utility, but one shouldn’t underestimate the difficulties. There are already numerous ‘secrecy’ laws that facilitate the construction and opacity of these systems, and these would have to be dismantled at

national and supranational levels. Furthermore, the technologies that allow for this are already woven into our everyday lives and so it makes sense for similarly structured governance systems to be built on top of them. There are also powerful ideological forces behind their creation and use.[48] Resistance and regulation may be futile in this domain. Could biomedical enhancement help? There are reasons to think that it could. First, biomedical enhancements could help to improve the capacities and efficiencies of human workers within governance systems, thereby countermanding some of the drive towards their replacement. Imagine that instead of relying on an automated facial recognition algorithm, the registry of motor vehicles encouraged its employees to adopt neural prosthetics that could augment their visual acuity and speed, and enable them to see the suspiciously similar matches. Second, biomedical enhancements that are more widely distributed could help to improve the capacities of all people affected by those governance systems. For instance, widespread cognitive enhancement, in conjunction with educational programmes, could enable people to better understand and participate in an increasingly automated governance system. The pharmacological enhancements could focus on core capacities and competencies – concentration, memory retrieval, basic reasoning capacities – and the education programme could focus on explaining the mechanics of automated governance. This would help to combat some of the proceduralist concerns raised in the previous paragraphs. These biomedical enhancements could work in tandem with regulatory interventions that give people the right to challenge and access these systems, thereby doing a better job conserving this part of the current evaluative equilibrium. So the claim here is not that biomedical enhancements solve the problem on their own, but rather that they could be an essential piece of the puzzle: without enhanced reasoning capacities people will be overwhelmed and poorly disposed towards acquiring the requisite knowledge that will allow them to conserve the value of participative democratic governance.

The final threat arises in the personal sphere, *i.e.* the day-to-day activities that make up our lives. There is a general desire to live a ‘good life’, where this can be understood as a life of meaning and virtue. Work plays a part in this conception of the good life but obviously work itself is not the be-all and end-all of the well-lived life. We also seek meaning and virtue in our family lives, leisure pursuits, unpaid work, charitable endeavours, sporting activities and so on. In these pursuits, a sense of personal achievement is often seen as central to whether they provide meaning or not.[49, 50] The longing for achievement is part of the current evaluative equilibrium:

achievement is a source of value and something valued by us in and of itself. But achievement is only really made possible when there is an appropriate causal-mental link between what an individual does and the outcome they are trying to achieve. The problem is that the rise of AI and robotics threatens to sever that link. This happens through the rise of personalised automation. It is already obvious that individuals rely on assistive software technologies in their everyday lives, for instance, software mapping technologies and personal digital assistants. There is also the emerging field of personal robotics, which includes robots that mow the lawn, vacuum the floors, cook your meals and so on. Reliance on such technologies can be expected to grow. Not only does this threaten to sever the link necessary for personal achievement, it also gives rise to a disenchanting effect on personal cognitive capacity.[51] Several studies demonstrate that the use of assistive software can lead to the degeneration of the cognitive capacities needed to perform certain mental and physical tasks.[52-54] This can occur even if there is a residual causal-mental link between what the individual does and the outcome in the world, provided that residual link is mediated through some degenerating technology.

This threat to the evaluative equilibrium may be less severe than the previous two. Personalised automation has its merits and demerits; it all depends on how it is used. Some people could use it wisely to outsource mundane or limiting activities and spend the rest of their time on activities that really do contribute to the good life. But we should be cautious in our optimism. People can be lazy and prone to biases that cause them to do things that are counterproductive to their own well-being. Overreliance on personalised automation could exacerbate these tendencies. This will be particularly true if the technology promises to do a better job of achieving the relevant outcome than a human ever could or if it supplies addictive hedonic rewards that distract us from those achievements.⁷ Biomedical enhancement can, once again, be part of the conservative solution to these problems. First, mood enhancers and cognitive enhancers could be adopted to mitigate or limit the biases that drive overreliance on personalised automation. And second, it could be used to enhance human capacities so that there is no need to turn to those technologies for better outcomes.

⁷ An extreme example would be the vision of the future depicted in the Pixar movie Wall-E. The imagined future consists of innumerable robots that perform routine tasks and supply humans with all their needs and pleasures. The humans are reduced to passive, overweight beings, moving about on automated chairs, receiving a constant and addictive supply of sugary foods and light entertainment. I suspect this possible future is highly unlikely to materialize, but its exaggeration of the human susceptibility to short term rewards is instructive.

In summary: work, legitimate governance and personal achievement are prominent features of our current evaluative equilibrium; developments in AI, robotics and information technology threaten those features; but biomedical enhancements could be used to conserve them. There is empirical uncertainty within this argument but wariness in the face of an uncertain future is part and parcel of the conservative position favoured in this article. The argument makes a strong claim about the conservative power of biomedical enhancement. The features of our evaluative equilibrium that have been singled out by the argument are not mere sideshows. They are central to our current economic, political and social institutions. Furthermore, the threats posed by the automating technologies are not insignificant. The erosion of work, comprehensible governance, or personal achievement individually constitute significant disruptions to our evaluative equilibrium. When two or more of them combine, the effect is truly radical. The evaluative conservative should be strongly committed to minimising such disruption and biomedical enhancement may, according to the arguments proffered above, be the most viable form of conservation.

5. Objections and Replies

I close by considering six objections to this argument.

The first objection is: why can't the conservative adopt a strategy of general resistance to technological development? My argument supposes that they must accept the reality of workplace automation, algorithmic outsourcing in government, and increased infiltration of AI and robotic assistance into everyday life. Surely it is possible and preferable to resist those developments *and* resist biomedical enhancement as well?

There is some merit to this criticism. It probably would be ideal, from the conservative's perspective, to resist all the developments since that would conserve more of the existing evaluative equilibrium. The question is whether this is feasible. The conservative may have to pick and choose their battles. When discussing the rise of algorithmic governance in the previous section, I suggested that it might be difficult to resist this phenomenon due to the fact that: (i) the underlying technologies are woven into the fabric of everyday life; (ii) there are laws and policies in place that facilitate and encourage their development; and (iii) there are powerful ideological forces driving

their adoption. The same holds in the case of workplace automation. Indeed, automation in work is deeply engrained into the incentive structure of capitalism. It is less true in the personal sphere but automation is becoming increasingly normalised there too and many of same forces underlying the drive towards workplace automation are behind the push towards personal automation. The question then becomes whether it is possible to resist the push towards automation across all three domains. It may be, but one advantage of biomedical enhancement is that it can facilitate conservation of values across all three domains at the same time. This is because success for humans across all three domains is dependent on general human intelligence and ability, and biomedical enhancement tends to target generic cognitive, emotional and perceptual capacities that are relevant to value conservation across all three domains. Until we develop generalized machine intelligence, the technologies that pose a threat to all three domains will be diverse and disparate.

The second objection is that even if biomedical enhancements preserve the values mentioned above, couldn't it also be the case that they change other values that the conservative favours? Sandel's giftedness argument illustrates the objection.[55] He argues that one prominent and important feature of our existing evaluative equilibrium is our gratitude for the gifted aspects of our lives. That is to say: gratitude for the fact that many good things happen to us because of biological and sociological factors that are beyond our control. This has a number of positive consequences in terms of empathy and social solidarity.[14] He worries that if we continue to pursue biomedical enhancement we will lose this element of our evaluative equilibrium because everything will become contingent and manipulable. Similarly, there are those who worry that biomedically enhanced personal achievement is less valuable because it is less authentic.

Three things can be said in response to this. First, these forms of conservative argument are often unpersuasive on their own merits. 'Authenticity' can be parsed in different ways, sometimes as a fear about how biomedical enhancements bypass personal effort, other times as a fear about disconnection from the authentic self. But there is no reason why biomedical enhancements have to eliminate effort: augmenting a capacity does not mean you stop using it in an effortful way.[56] Nor is it clear why it would lead to disconnection from the authentic self: if the decision to enhance is undertaken with appropriate self-reflection then it can be an expression of the authentic

self. Second, biomedical enhancements could be prioritized and selected on the basis that they help to conserve the values identified in these arguments. For instance, feelings of gratitude, empathy and solidarity are things that can be conserved and enhanced through various pharmacological and behavioural interventions.[62] Finally, even if this is all wrong, part of the motivation behind the argument in this article is that the evaluative conservative may sometimes face important tradeoffs when it comes to the conservation of value. The question then becomes whether the threat to something like gratitude and authenticity outweighs the three threats I have identified.

A third objection holds that there is a paradox at the heart of my argument. My claim is that externalised automation technologies pose a series of threats to our current evaluative equilibrium. This is because these technologies take over the performance of tasks traditionally performed by humans. But would things really be any better if (enhanced) humans continued to perform those tasks? It is well-known that the routine performance of a task by a human gives rise to *internal automaticity*, *i.e.* the task is initially learned and performed through conscious mental representations, but as the human gets better it transfers to subconscious systems within the human body. This is particularly true for many tasks performed at work, as well as certain hobbies and skills. If everything ends up automated in the end (either through external automation or internal automaticity) what difference does it make whether it is mediated through the human body?

Two responses seem apposite. First, this objection has a limited scope. While it is true that many human tasks end up being internally automatised, it is not always true. Certain cognitive tasks are always performed through conscious mental representations and their continued conscious representation may be part of what makes them valuable. For instance, the conscious understanding of political and bureaucratic decision-making may be essential to their legitimacy -- nothing less will do. Second, the conservative should be comfortable with the notion that there is something preferable about the current system (even we cannot pinpoint exactly what it is), whereby tasks end up automatised through the mediation of the human mind, as compared to a potential system whereby the automation is fully external.

The fourth objection is similar and makes use of the extended mind hypothesis (EMH) to argue that external and internal enhancements are ethically on a par.[57-59]

As Levy puts it: ‘alterations of external props used for thinking are (*ceteris paribus*) ethically on a par with alterations of the brain’.⁸ Adopting this, one might argue there is no principled difference between enhancements internal to the human body and externalised automating technologies, provided the combined internal and external system reaches the same outcome.

But the ‘*ceteris paribus*’ clause in this principle is all important. There does seem to be a qualitative difference between, say, a calculation performed in the mind and with the help of a calculator. The former seems to be more of a personal achievement than the latter. This points to a general problem with the use of the EMH in this debate. Even if you are persuaded that a functionalist account of the mind lends support to the EMH, you can still acknowledge that there are sufficient, relevant differences between mental processes realised internally (i.e. using biological systems) and those that are realised using a combination of internal and external mechanisms. These differences can include the phenomenology of the underlying mental activity, as well as how the activity gets dynamically integrated into the cognitive framework the individual uses to perceive and respond to the world.[60, 61] These differences are grist for the conservative mill. We should hold onto values we know attach to predominantly internalised activity, rather than the (potentially different) values that attach to predominantly externalised ones.

The fifth objection tries to argue that my conservative argument is flawed because it ignores the risks and impracticalities that arise from biomedical forms of enhancement. Biomedical enhancement involves interventions that operate on biological systems. But the human body is a complex, evolved system. Tinkering with it through pharmacological or prosthetic enhancements could be risky. Agar puts the point nicely by noting that all proponents of internal biomedical enhancement face an *integration problem*,[13] whereas proponents of externalisation can ‘cleverly avail themselves of efficiencies enabled by biological design’ (p. 47). This seems to make externalisation the conservatively superior option.

There are several things to be said in response. First, when Agar describes the integration problem he assumes that the externalisation has an enhancing effect on individual humans. I have already argued that this may not be true. Widespread

⁸ Levy [59] quoted in DeMarco and Ford [60] at 318. This is the strong version of the EPP. A weaker version states that they are equivalent insofar as the same ethical reasons apply to internal and external forms of enhancement.

externalisation could have severe disenchanting effects by making individuals more likely to be worse off relative to the current status quo (e.g. through unemployment) or through the atrophying of cognitive capacities. There are, no doubt, risks inherent in the use of biomedical enhancements that do not arise in the case of externalisation. But that doesn't make externalisation the preferable option. The whole point of the argument I have just defended is that externalisation poses significant risks to our current evaluative equilibrium and it is difficult to mitigate those risks without appealing to some degree of biomedical enhancement. Furthermore, the argument I am making is not that biomedical enhancement should be pursued at all costs. It is simply that biomedical enhancement should be prioritised given its conservative potential.

The final objection is that my argument is unrealistic due to the competing timelines for technological progress. Development of automating technologies is rapid and possibly exponential; development of biomedical enhancements seems to be much slower. We have several putative biomedical enhancers available today, but their effects are rather modest. We could pump more money into their research and development, but there is no guarantee of success. External automating technologies may win the race and we may lose much of our current evaluative equilibrium regardless of our efforts.

My response to this is simple: this may indeed be true, but it only serves to strengthen the argument I am making. If the risk to our evaluative equilibrium really is this great, then evaluative conservatives should be even more wedded to prioritising the proposed solution.⁹

5. Conclusion

In conclusion, despite the widespread belief that conservative moral principles are opposed to human enhancement, there are in fact strong reasons to think that human enhancement has conservative potential. This is because technological development

⁹ A reviewer worries that the argument implies that individuals – through their decisions to enhance or not – rather than collective institutions bear responsibility for conserving our evaluative equilibrium. They noted that this could be problematic if individuals lacked the means to avail of such interventions or were driven by their social context to avail of them in problematic ways (e.g. the poor worker who uses biomedical enhancements to resist bad working conditions but ignores side effects). But this is not a necessary implication of the argument. Although the ultimate *target* of enhancement is the individual, there is no reason why the enhancements themselves cannot be developed and distributed through collective institutions, nor why the individual must bear responsibility for the failure to enhance, or for the negative side effects of enhancement. For a longer discussion see Danaher [62]

does not take place in a vacuum. One cannot consider the effects of biomedical enhancement technology in isolation from other trends in technological progress. When this is done, it becomes apparent that AI, robotics and information technology are developing at a rapid pace and their widespread deployment could undermine much of our current evaluative equilibrium. Biomedical enhancement may be necessary, not merely desirable, if we are to maintain that equilibrium.

Acknowledgments: The author would like to thank Brian Earp and three anonymous reviewers for feedback on earlier drafts of this article.

References

1. Cohen, Gerald Allan (2011). Rescuing Conservatism. In *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*, eds. Wallace, R. Jay, Rahul Kumar, and Samuel Freeman. Oxford: OUP.
2. Cohen, Gerald Allan (2012). Rescuing Conservatism: A Defense of Existing Value. In *Finding Oneself in the Other*, ed. Otsuka, Kazuhiro. Princeton: Princeton University Press.
3. Pugh, J., Kahane, G. and Savulescu, J. (2013). Cohen's Conservatism and Human Enhancement. *J Ethics* 17(4): 331 – 354.
4. O'Hara, K. (2011). *Conservatism*. London: Reaktion Books.
5. Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. Oxford: OUP.
6. Guerrero, A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philos Stud* 136: 59-97
7. Moller, D. (2011). Abortion and Moral Risk. *Philosophy* 86: 425-443
8. Sepielli, A. (2008). What to Do When You Don't Know What to Do. In Shafer-Landau (ed) *Oxford Studies in Metaethics*. Oxford: OUP.
9. Buchanan, A. (2011) *Beyond Humanity? The Ethics of Biomedical Enhancement*. Oxford: OUP.
10. Kass, L. (2003). Ageless Bodies, Happy Souls: Biotechnology and the Pursuit of Perfection. *The New Atlantis*.
11. Fukuyama, F. (2002). *Our Posthuman Future*. New York: Farrar, Strauss and Giroux.
12. Agar, N. (2011). *The End of Humanity*. Cambridge, MA: MIT Press
13. Agar, N (2013). *Truly Human Enhancement*. Cambridge, MA: MIT Press
14. Hauskeller M (2011) Human enhancement and the giftedness of life. *Philosophical Papers* 40(1): 55–79.

15. Hauskeller M (2013) *Better humans? Understanding the enhancement project*. London: Acumen.
16. Savulescu J, Sandberg A, Kahane G. Well-Being and Enhancement (2011). In: Savulescu J, Meulen R ter, Kahane G, editors. *Enhancing Human Capacities*. Oxford: Blackwell Publishing Ltd.
17. Kahane G, Savulescu J. (2009). The welfarist account of disability. In: Brownlee K, Cureton A, editors. *Disability and Disadvantage*. Oxford: Oxford University Press p. 15–53
18. Juengst E. (1998). What does enhancement mean? In: Parens E (ed) *Enhancing Human Traits: Ethical and Social Implications*. Georgetown: Georgetown University Press.
19. Earp, B, Sandberg, A, Kahane, G and Savulescu, J (2014). When is diminishment a form of enhancement? Rethinking the enhancement debate in biomedical ethics. *Front Syst Neurosci* 8:12
20. Daniels N. (2000). Normal Functioning and the Treatment-Enhancement Distinction. *Cambridge Q Healthc Ethics* 9:309–22
21. Harris, J. (2007) *Enhancing Evolution: The Ethical Case for Making Better People*. Princeton NJ: Princeton University Press.
22. Parens, E. (1998). Is Better Always Good? In Parens (ed) *Enhancing Human Traits*. Washington, DC: Georgetown University Press.
23. Loi, Michele (2015). Technological Unemployment and Human Disenhancement. *Ethics Inf Technol* 17(3): 201-210.
24. Kurzweil, R. (2006). *The Singularity is Near: When Humans Transcend Biology*. London: Penguin
25. Agar, N. (2015) *The Sceptical Optimist*. Oxford: OUP.
26. Pratt, G. (2015). Is a Cambrian Explosion Coming for Robotics. *J Econ Perspect* 29(3): 51-60.
27. Ford, M. (2015) *The Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic Books.
28. Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton and Co.
29. Kaplan, J. (2015). *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence*. Yale University Press.
30. Susskind, R and Susskind, D. (2015) *The Future of the Professions*. Oxford: OUP.
31. Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms and that Control Money and Information*. Harvard: Harvard University Press.
32. Dormehl, L (2014). *The Formula: How Algorithms Solve All our Problems and Create More*. Perigee Books.
33. Lupton, D (2016). *The Quantified Self*. Wiley-Blackwell.

34. Gheaus, A., & Herzog, L. M. (2016). The goods of work (other than money). *J Soc Philos*. Forthcoming, available at <http://philpapers.org/rec/GHETGO>. Accessed 13 March 2016.
35. Autor, D. (2015). Why are there still so many jobs? The history and future of workplace automation. *J Econ Perspect* 29(3): 3-30
36. Srnicek, N and Williams, A (2015). *Inventing the Future: Postcapitalism and a World Without Work*. London: Verso.
37. Hughes, J (2014). A Strategic Opening for a Basic Income Guarantee in the Global Crisis Being Created by AI, Robots, Desktop Manufacturing and BioMedicine. *J Evol Tech*. 24(1): 45-61.
38. Gaus, G. (2010). *The Order of Public Reason*. Cambridge: Cambridge University Press.
39. Peter, F. (2014) Political Legitimacy. In Edward N. Zalta (ed) *The Stanford Encyclopedia of Philosophy Spring 2014 Edition* -- available at <http://plato.stanford.edu/archives/spr2014/entries/legitimacy/>
40. Aneesh, A. 2006. *Virtual Migration*. Duke University Press.
41. Aneesh, A. 2009. Global Labor: Algocratic Modes of Organization. *Sociological Theory* 27(4): 347-370
42. Danaher, J. 2016. The Threat of Algocracy: Reality, Resistance and Accommodation. *Philos Technol* DOI: 10.1007/s13347-015-0211-1
43. Zarsky, T. (2013). Transparent prediction. *University of Illinois Law Review* 4: 1504.
44. Zarsky, T. (2015). The Trouble with Algorithmic Decision-Making: An Analytic Road Map to Examine Efficiency and Fairness in Automated and Opaque Decision Making' (2015) *Sci Technol Human Values* DOI: 10.1177/0162243915605575
45. Burrell, J (2016). How the machine 'thinks': understanding opacity in machine learning algorithms. *Big Data Soc* DOI: 10.1177/2053951715622512.
46. Kitchin, R. (2016). Thinking critically about and researching algorithms. *Information, Communication and Science*. DOI: 10.1080/1369118X.2016.1154087
47. Estlund, D. (2003). Why not Epistocracy? In Naomi Reshotko (ed) *Desire, Identity, and Existence: Essays in Honour of T.M. Penner*. Academic Printing and Publishing.
48. Zuboff, S. (2015). Big Other: Surveillance Capitalism and the Prospects of an Information Civilization. *Journal of Information Technology* 30: 75-89
49. Luper, S. (2014). Life's Meaning. In Luper, S (Eds) *The Cambridge Companion to Life and Death*. Cambridge University Press.
50. Wolf, Susan. 2010. *Meaning in life and Why it matters*. Princeton, NJ: Princeton University Press.
51. Carr, N (2015) *The Glass Cage: Where Automation is Taking Us*. New York: WW Norton & Co.
52. Van Nimwegen, C. et al (2006). The Paradox of the Assisted User: Guidance can be Counterproductive. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* 917-926

53. Burgos, D. et al (2007). Game-based learning and the role of feedback. A case study. *The International Journal of Advanced Technology in Learning* 4(4): 188-193.
54. Dowling, C, Leech, SA and Moroney, R (2008). Audit Support System Design and the Declarative Knowledge of Long-Term Users. *Journal of Emerging Technologies in Accounting* 5(1): 99-108.
55. Sandel M (2009) *The case against perfection*. Harvard: Belknap.
56. Douglas T (2014) Enhancing moral conformity and enhancing moral worth. *Neuroethics* 7(1):75-91
57. Levy, N. (2007a). *Neuroethics: Challenges for the 21st Century*. Cambridge: Cambridge University Press.
58. Levy, N. (2007b). Rethinking Neuroethics in Light of the Extended Mind Thesis. *Am J Bioeth* 7(9): 3-11.
59. Levy, N. (2011). Neuroethics and the Extended Mind. In Sahakian, B. and Illes, J (Eds) *Oxford Handbook of Neuroethics*. Oxford: OUP.
60. DeMarco, J and Ford, P. (2014). Neuroethics and the Ethical Parity Principle. *Neuroethics* 7: 317-325
61. Michaelian, K (2012). Is External Memory Memory? Biological Memory and the Extended Mind. *Consciousness and Cognition* 21(3): 1154-1165.
62. Danaher, J. (2016). Human Enhancement, Social Solidarity and the Distribution of Responsibility. *Ethical Theory Moral Pract* 19(2): 359-378.
63. Persson, I and Savulescu, J. (2012). *Unfit for the Future*. Oxford: OUP.