

order element is not a part of the conscious state, but is extrinsic to it. And this opens them up to worries about misrepresentation. On extrinsic views, we can get a radical mismatch between our awareness of the conscious state and the conscious state itself. What would it be like for the subject in such cases? Objectors like Alex Byrne [1997] and Joseph Levine [2001] contend that whatever the higher-order theorist says, the view breaks down into incoherence. I will not here speak on the cogency of this worry, but Gennaro accepts the objection and builds his view to avoid the difficulty altogether. If the higher-order part sufficiently misrepresents the lower-order element, then on the WIV, there is no conscious state at all. The entire complex *in good working order* is necessary for experience. So the worry does not arise.

However, because Gennaro invokes a broadly causal-functional view of representation (indeed, this is what allows for the naturalistic reduction of the WIV), it seems that one could in principle have badly mismatched parts. There is no causal prohibition on such a possibility. And then we can ask what it would be like for the subject in such a case. If the experience is simply what the higher-order element dictates, the lower-order world-representing element falls away as otiose. The higher-order thought does all the phenomenal work and the lower-order sensory representation—the red and yellow sensations of a sunset, the rich taste and aroma of good coffee—don't matter for what it's like for the subject. And this is seen as a disastrous collapse of the view. But if one holds that the lower-order element fixes what it's like for the subject, we lose a role for the higher-order awareness. Further, we are left wondering why these lower-order elements are conscious. The higher-order state allegedly making them conscious can shift its content radically without altering how things appear to the subject. We've lost the relational explanation of experience distinctive of the higher-order approach. Gennaro can protest that a special kind of Kantian 'synthesis' is a necessary component of experience, and that this is lacking in the mismatch case. But now synthesis can't be given a straightforward causal-reductive analysis, and we are left wondering, paradoxically, how to explain consciousness.

While I am not convinced that the danger of misrepresentation is fatal to the higher-order view, I don't see that Gennaro has avoided it and this leaves his view on all fours with traditional higher-order theories. But in the final analysis, *The Consciousness Paradox* is a spirited defence of the higher-order approach and the WIV. Gennaro may not have fully dispelled the air of paradox, but he has made clear that the approach is a viable and important contender in the effort to crack the hard problem of consciousness.

Josh Weisberg

*University of Houston*

© 2013 Josh Weisberg

<http://dx.doi.org/10.1080/00048402.2013.831930>

Persson, Ingmar, and Julian Savulescu, *Unfit for the Future? The Need for Moral Enhancement*, Oxford: Oxford University Press, 2012, pp. ix + 143, £21.00 (hardback).

Postulate a good great enough or a crisis dire enough and any possibility—no matter how far-fetched—will bear investigation. This has proved a productive rhetorical strategy in applied ethics in recent years, enabling philosophers to pontificate in print about the dangers of nanotechnological 'grey goo', the need to ensure that our robot masters will be friendly after the coming Singularity, and the ethics of genetically modifying people to mitigate the effects of climate change—to name just a few

projects—despite the near-complete absence of our capacity to bring these any of these things about. Arguably, the latest example of this phenomenon is Ingmar Persson and Julian Savulescu's book *Unfit for the Future*, which advertises the need for what the authors call 'moral bioenhancement' in order to avoid what they describe as 'ultimate harm'. The book is essential reading for those wanting to understand the flourishing debate about 'moral enhancement' that has emerged in bioethics in large part as a result of the authors' efforts in earlier publications. It also contains novel arguments in social and political philosophy. Finally, in so far as it represents the work of two of the leading figures in contemporary applied ethics functioning at the height of their powers, the book should be of wider interest to those working in this field.

Persson and Savulescu argue that humanity is threatened by a gap that has emerged between our technological powers and our moral psychology. We evolved in small tribal bands, where the existence of face-to-face relationships discouraged individuals from defecting from cooperative arrangements, and with a technology that placed significant limits on how much harm we could do to each other or the environment. Yet we now live in mass societies, in circumstances which make it easy for us to externalize the costs of our economic choices, and possess a scientific knowledge that grants us unprecedented power. The authors devote the second-longest chapter in this short book to detailing the 'common-sense morality' that is supported by our evolved psychological traits and argue that it is ill-suited to the conditions in which we now live. As a result of this mismatch, we have become a pressing danger to ourselves. In particular, humanity faces a grave threat from anthropogenic climate change and other environmental problems and from the catastrophic misuse of modern science, including the use of weapons of mass destruction. Persson and Savulescu suggest that in order to escape this predicament we require 'moral enhancement'. While they acknowledge the possibility of moral enhancement through education and social reform, the focus of the argument is on what they call 'moral *bio*enhancement', by which they mean the modification of individuals' moral psychology by the judicious application of pharmacology, neuroscience and, perhaps, genetic selection and/or modification.

Much of this argument has appeared in print elsewhere, so the book will be of most value to readers who are unfamiliar with the authors' other work, who wish to learn how the key claims fit together into a larger scheme, or who are interested to understand the evolution of the authors' thought. A significant difference between the book and the articles upon which it is based is that Persson and Savulescu are now explicit that moral enhancement alone will not solve the problems that concern them and that liberal societies will need to become significantly less liberal in order to address the 'threat' of weapons of mass destruction. Thus, they argue for a greatly increased use of surveillance technologies (which is hard to imagine in the context of recent revelations about the US government's activities in cyberspace) in order to protect society against the risk of 'ultimate harm'. This prospect troubles the authors less than it might others because the authors deny the existence of any right to privacy on the basis of a few brief thought experiments about the existence of people with super senses. More problematically, they also link the threat of terrorism to the presence of minorities with 'different political and religious ideals' in the midst of contemporary societies. Whilst they do admit that the possibility that fears of ultimate harm may be exaggerated and that concerns about terrorism may function in practice to reinforce xenophobia, it is striking how abstracted this discussion is from history or political science. A more interdisciplinary approach might, I suspect, have revealed that it is not our moral, but our political and institutional, failings that generate the threats with which the authors purport to be concerned.

Indeed, *Unfit for the Future* is marked throughout by a puzzling combination of optimism and pessimism.

Persson and Savulescu are profoundly pessimistic about the capacity of existing or realistically available political institutions to address the problems they highlight. Reviving Hobbes with a technological determinist twist, they insist that our current situation is dire because it is easier for people in modern societies to harm others than it is to help them. Yet, despite this, most people spend most of their time helping others through various cooperative activities, such as work, charity, and mutually beneficial social relationships rather than committing mass murder. Meanwhile, the authors' worry that 'terrorists' will acquire and use nuclear weapons has been a feature of political panic-mongering since the 1950s. Efforts to improve human security through political and institutional reforms, or to avert climate change by reducing our reliance on fossil fuels and/or abandoning economic models based upon the premise of never-ending economic growth are, they imply, doomed by the 'tragedy of the commons' and by our flawed moral psychology.

On the other hand, in order for the discussion of moral bioenhancement to be worth having at all, Persson and Savulescu must, at crucial points in the argument, be profoundly optimistic about the possibility of meaningful moral enhancement via technological means and, in particular, the possibility of influencing moral behaviour through pharmaceuticals. In order for moral bioenhancement to address the risk of ultimate harm they must be optimistic about the likelihood that any technology capable of modifying individuals' motivations and behaviour will be not be misused by authoritarian governments. One cannot help but worry that their argument that social problems are the result of individuals' being congenitally unfit for modern life and that the solution is to dose them up with (as yet imaginary) powerful mood-and-behaviour-altering drugs may in practice function mostly to license states' incarcerating the poor and the socially marginal and medicalizing deviancy.

Three important questions also receive surprisingly little attention in *Unfit for the Future*: what would it mean to say that someone was 'morally enhanced'; how would moral enhancement prevent the crisis they see as confronting humanity; and, how likely is it that we are going to be able use moral bioenhancement to achieve this?

The definition of moral enhancement is perhaps the most controversial question in the emerging literature on the topic. Persson and Savulescu opt for a relatively modest and plausible definition of moral enhancement as the modification of individuals' dispositions so as to make them more altruistic and to have more of a sense of justice. To establish the possibility of moral *bioenhancement*, they cite a small number of studies that show that drugs such as oxytocin and the selective serotonin reuptake inhibitors alter the extent to which people are inclined to cooperate with people in various rather artificial situations. These are interesting results but the general point that drugs can alter mood and behaviour will hardly come as news to anyone who has ever been drunk. Whether we should understand someone who donates to Oxfam because they were inebriated as having been 'morally enhanced' is a further—and important—question. Those who place the virtues or the possession of a 'good will' at the heart of their account of what it is to act morally, or who insist on egalitarian relations between persons as a condition of autonomy, are likely to be less enthusiastic about the prospects for moral bioenhancement than are the authors.

Nor is it clear that making individuals more altruistic and concerned for justice would necessarily solve the problems that exercise Persson and Savulescu. Reshaping our motives will not eliminate collective action problems, which are a product of the difference between what is individually rational and what is collectively rational, rather than of the motives of agents. Similarly, terrorism is a political phenomenon, stemming as much from an excess of moral sentiments as any deficiency thereof; in a

perverse fashion, suicide bombings are, it might be argued, the ultimate acts of altruism. Any attempt to rule out the possibility that 'enhanced' motivations will lead to outcomes that are no better than the status quo risks circularity in the definition of moral enhancement.

However, the real reason why moral bioenhancement seems unlikely to offer much by way of a solution to the threat of climate change and/or catastrophic misuses of modern science is that we don't have any such technology and we are unlikely to develop it within the time available to us to avert devastating climate change. Even if we did develop a technology that could magically reshape individuals' motivations as Persson and Savulescu desire, we would then have to solve the not-inconsiderable problems involved in applying it to 'hundreds of millions' [121] of people without presuming or licensing an authoritarianism that would most likely render the project moot.

At the end of *Unfit for the Future*, Persson and Savulescu admit that 'moral bioenhancement worthy of the name is practically impossible at present and might remain so for so long that we will not master it, nor succeed in applying it on a sufficient scale, in time to help us to deal with the catastrophic problems that we have outlined' [123]. Nevertheless, the authors suggest that 'it is a serious mistake to reject moral bioenhancement out of hand because the need for human moral enhancement is so acute that we should not write off any potential effective means without thorough examination' [121] and that 'it is important that moral bioenhancement is not written off without good reason' [123]. This latter claim at least seems eminently plausible, although it is significantly weaker than one might have expected from the opening chapters of the book.

Given the existing scientific consensus on the grave threat posed by anthropogenic climate change and the urgent need to reduce emissions of greenhouse gases now, talk of 'moral bioenhancement' as a potential solution is entirely fantastic. Of course, one can't help but suspect that the authors themselves do not really believe that moral bioenhancement offers any realistic prospect of avoiding climate change or reducing the risk of the use of weapons of mass destruction. Regardless, they have written a provocative defence of a bold thesis, which contains much to interest, challenge, and entertain their readers.

Robert Sparrow

Monash University

© 2013 Robert Sparrow

<http://dx.doi.org/10.1080/00048402.2013.860180>

Sterelny, Kim, *The Evolved Apprentice: How Evolution Made Humans Unique*, Cambridge, MA: MIT Press, 2012, pp. 264, US\$35 (hardback).

Kim Sterelny's *The Evolved Apprentice* is an extraordinarily rich and readable treatment of a range of problems associated with human evolution. Many of these problems are already well known, and here Sterelny treats the reader to novel solutions which are elegant and intuitive. But equally valuable are the lesser appreciated puzzles for which Sterelny provides both strong motivation to solve and a powerful framework for doing so.

In this review, I describe the core features of Sterelny's account, the general evolutionary and particular historical phenomena which it captures, and its virtues over some contemporary rival approaches. I will then offer remarks on some general features of the account that underlie its strength as an evolutionary model, followed by