

The biomedical enhancement of moral status

doi:10.1136/medethics-2012-101312

Russell Powell, Associate Editor

The biomedical enhancement of human capacities has emerged as one of the most philosophically invigorating areas of contemporary bioethical research. In exploring the ethical dimensions of emerging biotechnologies and human-machine interfaces, the literature on human enhancement has made significant contributions to traditional problems in moral philosophy. One such area concerns the enhancement of cognitive capacities that bear on moral status. Could biotechnological or other forms of neurocognitive intervention result in the creation of 'postpersons' who possess a moral status that is higher than that of 'mere persons'? If the creation of postpersons with a higher moral status is indeed possible, is it morally wrong to bring this about? Nicholas Agar takes up these two questions in his Feature Article for this issue (*see page 67*, Editor's Choice), in which he defends affirmative answers to both.

Most moral philosophers in the Kantian tradition would agree that were we to enhance the cognitive capacities of a non-rational sentient creature (such as a dog) so as to confer on it the psychological properties associated with personhood (such as practical rationality, self-awareness, interests that extend into the future, mutual accountability, and so on), we will have increased that creature's intrinsic moral worth relative to non-rational sentient beings, and thus we will have enhanced its moral status. Virtually all parties to the moral status enhancement debate agree on this much. Why then would it not be similarly possible in principle to enhance the cognitive capacities of mere persons so as to create postpersons who possess a higher moral status than that of mere persons? While this may seem like a reasonable extrapolation from the first scenario, it is here that moral philosophical boats begin to diverge.

Allen Buchanan was the first to address systematically the possibility of moral status enhancement¹ and his work sets the backdrop for Agar's analysis. Buchanan argues that personhood is a threshold rather than scalar property that, once realised, confers a single tier of moral status on its bearers. His chief basis for this assertion is our commitment to the equal moral worth of persons, which entails that having more or less of X (where X is some cognitive property or cluster of properties necessary for personhood) has no bearing whatsoever on moral status, so long as X is realised to some minimal (non-zero)

degree—and X is realised to some minimal degree, by definition, in all persons. The idea is powerful and seemingly persuasive: some people may be more rational, mutually accountable, deliberative, self-contemplative or forward-looking than others, but this does not imply that such individuals have greater moral worth or possess rights that are more inviolable, on any plausible account of human rights. In some cases it may be difficult to determine whether a minimal degree of X is present in a given animal, but this is a separate epistemic question.

Utilitarian moral philosophers, on the other hand, have an easier time justifying the differential weighting of mere person and postperson interests, respectively, according to the degree of X that they possess (for a discussion, *see page 80*). Agar's target, however, is the Kantian approach. He argues that our commitment to the equal moral worth of persons is contingent on the existing range of human variation in X, and fails to extrapolate to a scenario involving vast disparities in cognitive abilities, such as the mixed society of mere persons and postpersons that he envisions. Agar contends that not only is it possible, but that we actually have affirmative reasons to believe (see below), that there are higher moral status thresholds than the one occupied by mere persons—thresholds that can be realised through radical cognitive enhancement.

Yet a formidable challenge to the moral imagination remains: If there exists a higher threshold of moral status beyond that realised by mere persons, what sorts of cognitive capacities would give rise to this higher threshold? Would this simply involve greater degrees of the capacities canonically associated with personhood, or would it require, as suggested by Hauskeller (*see page 76*), an entirely novel cognitive ability, one that is different *in kind* from anything that mere persons possess? Agar punts on this question, though he has a good (if not entirely convincing) excuse: he contends that such a failure of imagination is to be expected, indeed is necessitated. For whatever property confers a higher moral status than that of mere persons, it will be a cognitive property; and because we, qua mere persons, lack that cognitive property, we cannot even conceive of what that property would be like, let alone describe its contours in any detail. Our position of epistemic ignorance vis-à-vis the defining properties of

postpersonhood, Agar analyses, is akin to the futility of a non-rational sentient creature attempting to contemplate cognitive abilities associated with a higher moral status, such as rationality. This move might be too quick, however. We human mere persons can recognise and remark on the ability of toothed whales and microchiropteran bats to form detailed images through echolocation, even though we manifestly lack this ability. In a similar vein, why couldn't we imagine nomicallly possible cognitive properties that we lack, such as visualising the universe in higher dimensions? Of course, there is no reason to think that abilities such as these have anything to do with moral status. But that is precisely the point. Our failure to imagine what properties might confer higher moral status on postpersons, if they do not simply involve more of X, is neither expected nor nomicallly necessitated.

Nevertheless, Agar offers an inductive argument that such morally relevant properties exist. He notes that we currently have at least three widely accepted levels of moral status: non-sentient entities (with zero moral status), sentient non-persons (with intermediate moral status), and persons (with the currently highest level of moral status). We might add to this list (though not without controversy): some environmental ethicists hold that non-sentient living things possess a non-zero moral status, and many bioethicists have argued for a minimal moral status in the case of human embryos. Agar's induction proceeds as follows: given the plurality of putative moral statuses in the known universe, it seems likely that there are higher levels of moral status that are possible but currently not instantiated, or actual but currently unknown. To think otherwise, if I may offer an analogy on Agar's behalf, would be akin to tallying the number of currently observable exoplanets and claiming, quite ludicrously, that there are precisely no more to be found.

Note, however, the following: if Agar is right that we cannot conceive of the cognitive property that confers higher moral status, then it is impossible to know what sorts of cognitive enhancements would bring about higher moral status—the implication being that moral status enhancement could only occur as an accidental or incidental side effect of the enhancement of known, targeted properties of mere persons. If this is so, then Agar should be wary of even moderate

cognitive enhancements, since their unexpected side effect could be to create properties that confer higher moral status.

Having inferred that radical cognitive enhancement could create beings with higher moral status, Agar pushes on to argue that we have moral reasons to refrain from creating such beings. His worry is not merely that postpersons could wrong mere persons by, say, enslaving or exploiting them, or otherwise treating mere persons in ways that are inconsistent with their actual moral status. This is surely a grave risk, one that Agar accords insufficient weight to in his article, given the deep and dark history of human rights violations between mere persons, not to mention the countless moral atrocities inflicted by mere persons on non-rational sentient beings (many of which are ongoing). Agar's primary concern, however, is that postpersons may have legitimate interests the satisfaction of which requires or permits the *morally justified* sacrifice of mere persons.

Few believe, with Kant, that the inviolability of persons is absolute. But if it is permissible to sacrifice mere persons in the case of 'supreme emergencies', as Jeff McMahan has suggested, then it is possible that postpersons will have greater immunity than mere persons in such cases.² Agar is worried not so much with the sacrifice of mere persons in the context of supreme emergencies (which he thinks may be rare), but in relation to what he calls 'supreme opportunities'—situations that promise great benefits that would accrue to postpersons by virtue of sacrificing mere persons. Agar thinks that supreme opportunities will commonly present, and he infers that such sacrifices will be morally permissible when beings of different moral status are involved. His inference rests on the observation that we find it morally acceptable to carry out painful and deadly experiments on non-human primates (but not human persons) in order to find cures for serious human diseases. This inference might be unwarranted, though, for two reasons.

First, the fact that mere persons commonly sacrifice non-rational (and lesser-rational) sentient animals for a variety of human purposes without moral consternation is not particularly probative in itself, since many practicing experimentalists, as well as members of the public and the academy who support their research, frequently serve up indefensible speciesist rationales for these practices. In actual fact, philosophers are divided on the issue of painful primate experimentation even in the context of meaningful clinical research. Sacrificing non-rational sentient beings in the context of supreme opportunities is not as widely embraced as Agar seems to suggest (especially once we take into account our intuitions about non-rational sentient human beings), and thus the moral permissibility of such sacrifices

is very much in question. Douglas (*see page 75*) offers a similar critique of Agar's induction, stating that 'the relevant cases from which to extrapolate is not the actual world, but one in which persons exist and *treat animals only in morally permissible ways*' (emphasis in the original).

Secondly, the observation that sentient non-persons are sacrificed to procure benefits for mere persons, assuming these sacrifices are morally justifiable, is consistent with Buchanan's strong moral threshold theory and Agar's multiple moral threshold theory. On Buchanan's view, it is possible (though by no means a given) that sentient non-persons could be permissibly sacrificed to procure significant benefits for persons because the latter possess attributes associated with personhood, which confer full moral status, while the former possess a truncated set of rights (eg, the right to not be made to suffer needlessly). Since the thing observed (non-human primate experimentation) may be permissible on each rival theory of moral status, it does not increase the probability that the multithreshold account is correct, and hence it cannot serve as a basis by which to adjudicate between them. Wasserman (*see page 78*) argues in favour of the opposite induction, given that the trajectory of moral progress from Ancient times to the contemporary human rights regime has involved a consistent *reduction* rather than multiplication of moral status distinctions. Recalling the discussion above, it is as if the number of observed exoplanets had been shrinking.

Agar concludes that we ought not to bring postpersons into the world, because (1) we have no moral obligation to bring postpersons into the world, and (2) there is a reasonable probability that doing so would entail undesirable—indeed, tragic—consequences for mere persons. Agar argues for the second proposition, but he assumes the first, to the peril of his argument. Some philosophers (including the Editor-in-Chief of this journal) have proposed that parents have a duty to create offspring who are expected to have the best life—and a plausible (though by no means slam dunk) case can be made that postpersons can be expected (on average) to live a better life than individuals with the cognitive capacities of mere persons, insofar as postpersons are capable of enjoying goods and engaging in levels of cooperation and flourishing that are inaccessible to mere persons. It is not immediately clear why the above duty should be limited to parents, as similar consequentialist considerations would seem to give rise to a more general duty to create individuals and social states of affairs associated with higher levels of wellbeing. Douglas argues along these lines when he suggests that the creation of postpersons would be a valuable

addition to the world, just as the addition of sentient persons is (*ceteris paribus*) a valuable addition to a world comprised solely of sentient non-persons.

Thus, for Agar's argument to go through, he would have to resist the consequentialist obligation to bring about valuable states of affairs, deny that societies with postpersons are more valuable than societies without them, or accept that these considerations give rise to a *prima facie* obligation to create postpersons but argue that it is outweighed by the potential harms to mere persons or the intrinsic value of remaining merely human. But such an all-things-considered analysis would vitiate the claim that there is no obligation to create postpersons. Persson (*see page 77*) takes up a similar thread, contending that Agar's analysis is biased against postpersons in that it considers only the potential harm to mere persons. Moreover, the creation of postpersons could, mere Persson argues, help mere persons to avoid impending military and environmental catastrophes. Agar's argument would thus be on stronger logical footing, albeit stripped somewhat of its urgency, were it couched in terms of a *prima facie* obligation to refrain from creating postpersons. For at the end of the day it is not clear that the risk of justified harms to mere persons outweighs the benefits redounding to the same by virtue of having postpersons around.

Finally, because Agar's concerns arise from envisaged synchronic cognitive discrepancies between postpersons and mere persons, they would dissipate if a society of postpersons were achieved in a gradual, piecemeal fashion, such that the variance in cognitive capacity at any given time slice of the moral population was never great enough to result in the occupation of different moral status thresholds above that of mere persons. If we have moral reasons to create a world of postpersons, we may be able to achieve this end, while avoiding the costs that Agar fears, by engaging in moderate enhancements that lead us on a steady but sure path to a society of postpersons. The moral resources to oppose this result are not, so far as I can tell, present in Agar's essay.

In his reply to commentaries, Agar addresses many of these criticisms, and I leave it to the reader to judge whether he does so successfully (*see page 81*). The enhancement of moral status remains a ripe topic for theoretical investigation, with philosophical payoffs redounding well beyond the field of bioethics. It is my hope that discussions of this calibre will continue to grace the pages of this journal.

REFERENCES

1. Buchanan AE. *Beyond Humanity?* Oxford: OUP Oxford, 2011:209-42.
2. McMahan J. Cognitive disability and cognitive enhancement. *Metaphilos* 2009;40:582-605.