

## Human enhancement and supra-personal moral status

Thomas Douglas

Published online: 26 August 2011

© The Author(s) 2011. This article is published with open access at Springerlink.com

**Abstract** Several authors have speculated that (1) the pharmaceutical, genetic or other technological enhancement of human mental capacities could result in the creation of beings with greater moral status than persons, and (2) the creation of such beings would harm ordinary, unenhanced humans, perhaps by reducing their immunity to permissible harm. These claims have been taken to ground moral objections to the unrestrained pursuit of human enhancement. In recent work, Allen Buchanan responds to these objections by questioning both (1) and (2). I argue that Buchanan’s response fails. However, I then outline an alternative response. This response starts from the thought that, though moral status-increasing human enhancements might render ordinary, unenhanced humans less immune to permissible harm, they need not worsen the overall distribution of this immunity across beings. In the course of the argument I explore the relation between mental capacity and moral status and between moral status and immunity to permissible harm.

**Keywords** Moral status · Personal identity · Inviolability · Harm · Partiality · Cognitive enhancement · Moral enhancement · Human enhancement

A typical adult human enjoys a special moral status. This status is often thought to confer certain basic rights or claims—for example, to self-determination and freedom from some forms of non-consensual interference. I will use the term ‘persons’ to refer to the class of beings with the moral status (or one of the moral statuses) characteristic of currently typical adult humans.

---

T. Douglas (✉)  
Balliol College, University of Oxford, Broad Street, Oxford OX1 3BJ, UK  
e-mail: thomas.douglas@philosophy.ox.ac.uk

Many believe that our nearest primate relatives fail to qualify as persons; their moral status is lower than ours. This difference is often attributed to their lesser mental capacity. Perhaps it is due to their lacking rationality, practical rationality, or the capacity for moral agency. But if chimpanzees and other primates possess lower moral status than persons in virtue of their lesser mental capacity, we might speculate that beings with *greater* mental capacity than us would possess a higher moral status than persons—a supra-personal moral status. This possibility has long been a topic of theological speculation.<sup>1</sup> More recently it has been attributed practical significance in one of the liveliest debates in contemporary philosophical bioethics: it has been taken to ground an objection to the enhancement of certain human capacities.

## 1 The objection

Some drugs have been shown to enhance aspects of mental functioning in healthy individuals (de Jongh et al. 2008). To date, the demonstrated effects have been small. But further advances in neuropharmacology, brain–machine interface technologies and genetics may, in the future, enable the creation of beings whose mental capacity substantially exceeds our own, perhaps to a degree similar to that by which our capacity exceeds that of our nearest primate relatives. This has led some authors to speculate that

- (1) The technological enhancement of human mental capacities could result in the creation of beings with supra-personal moral status ('supra-persons').<sup>2</sup>

Arguably, mere persons could be permissibly harmed for the sake of these supra-persons in ways that they may not be permissibly harmed for the sake of one another. For example, perhaps persons could permissibly be used, without their consent, in medical experiments designed to aid supra-persons. Or perhaps persons could be rightly excluded from the democratic institutions of the supra-persons. This raises the concern that

- (2) The creation of supra-persons would harm ordinary, unenhanced humans.

Propositions (1) and (2) have been taken to support a range of moral reservations about human enhancements—interventions that augment the capacities of normal, healthy humans. These reservations have not been clearly specified, but in each case the initial thought seems to be that, at the very least

- (3) We (existing human persons) have moral reasons not to encourage certain human enhancements: those that would create supra-persons.<sup>3</sup>

<sup>1</sup> For example, in attempting to explain the origin of human suffering, the English writer and politician Soame Jenyns (1758, pp. 49–51) speculated that there existed “angels or demigods” with higher moral status than persons who could therefore rightly harm mere persons. For a satirical response, see Johnson (1825).

<sup>2</sup> The ‘could’ in this claim is to be read as denoting physical, not merely logical, possibility. The claim is that it is physically possible that technological enhancements will produce supra-persons.

<sup>3</sup> Moral reasons (henceforth sometimes merely ‘reasons’) are to be understood as *pro tanto* moral reasons: reasons with some, but not necessarily decisive, normative force. From (3), a range of further

Some protagonists in the ethical debate about human enhancement appear to have been moved by this line of thinking—from (1) and (2) to (3). In a critique of transhumanism, the movement most strongly committed to human enhancement, Fukuyama (2004, p. 42) writes that

Underlying [the] idea of the equality of rights is the belief that we all possess a human essence that dwarfs manifest differences in skin color, beauty, and even intelligence. This essence, and the view that individuals therefore have inherent value, is at the heart of political liberalism. But modifying that essence is the core of the transhumanist project. If we start transforming ourselves into something superior, what rights will these enhanced creatures claim, and what rights will they possess when compared to those left behind?

And elsewhere (2002, pp. 9–10),

The ultimate question raised by biotechnology is, What will happen... once we are able to, in effect, breed some people with saddles on their backs, and others with boots and spurs?

One of Fukuyama's worries is that enhanced beings might *claim* greater rights than are enjoyed by unenhanced humans. But another concern is that enhanced beings would *actually have* greater rights than the unenhanced. This could be true if they possessed supra-personal moral status, as (1) maintains they might.<sup>4</sup>

Even writers who are unpersuaded by most objections to human enhancement have taken the objection captured by (1)–(3) seriously. Allen Buchanan, Jeff McMahan and Julian Savulescu have all found it worthy of prolonged discussion,<sup>5</sup> and Buchanan (2011, p. 21) lists it among the eight most important concerns about the enhancement of human capacities. Buchanan's work provides the most sustained critique of the objection. In a recent article and book, he provides a rich and wide-ranging response to both (1), the suggestion that human enhancement might produce supra-persons, and (2), the claim that this would harm ordinary humans (Buchanan 2009, 2011). Though he does not claim to have decisively rebutted either proposition, he does cast considerable doubt on both. For example, he argues that "The idea of a moral status higher than that of persons is dubious, given a plausible understanding of the notion of moral status" and that "Even if we grant the dubious

---

Footnote 3 continued

conclusions might be inferred. It might be argued, for example, that, all things considered, we have most moral reason not to encourage enhancements that would create supra-persons, that we have most reason to take steps to prevent the development of technologies that would enable such enhancements, that our governments have most reason to ban the provision or use of such technologies, or at least that we and our governments may permissibly do these things.

<sup>4</sup> Others have worried that enhancement might render human rights obsolete. See, for example, Parens (1995). This concern could also be interpreted as an oblique expression of (1)–(3). Human rights could be rendered obsolete because they are eclipsed by broader or stronger rights that the enhanced beings have in virtue of their greater moral status.

<sup>5</sup> See Buchanan (2009, esp. pp. 346–369, 2011), McMahan (2009a, pp. 600–604), and Savulescu (2009). Some have also considered the possibility that mentally enhanced beings might have greater moral or legal rights than unenhanced ones, but not necessarily due to having greater moral status. See Buchanan (2009, pp. 371–381) and Wikler (2009).

assumption that the emergence of beings with a moral status higher than that of persons is possible, the emergence of [supra-persons] would not extinguish whatever rights the unenhanced have by virtue of being persons” (2009, pp. 349–350).

I will argue that (1) and (2) are in fact quite resilient—more resilient than Buchanan recognizes. Though we certainly lack conclusive evidence for their truth, we are not in a position to rule it out. We should remain open to the possibility that human enhancement could create supra-persons, and thereby harm ordinary, unenhanced humans. However, it is, I will argue, not clear that proponents of human enhancement need be troubled by this possibility, for there is an alternative way of responding to the objection captured by (1)–(3).

## 2 Preliminaries

Before embarking on my argument proper, I need to introduce some assumptions.

First, some assumptions about moral status. The concept of moral status is most commonly deployed in attempts to explain or justify the attribution of certain basic moral protections, such as basic moral rights or claims, to certain beings. For example, it may be said that human fetuses or non-human primates have a right not to be killed because they have the same moral status as ordinary adult humans. These attempts usually appeal, ultimately, to the non-moral properties of the beings in question, for example, their mental capacities, potential mental capacities, or species membership. The concept of moral status simply serves as a way of picking out those non-moral properties. To say that a being has a certain moral status is, on this view, roughly to say that it has whatever intrinsic non-moral properties give rise to certain basic moral protections.

In accordance with this way of thinking about moral status, I will assume that

**Assumption I** Other things being equal, a being with higher moral status will enjoy stronger and/or broader basic rights or claims than a being of lesser moral status.

Beyond this, I wish to remain as open as possible to competing views about the relationship between moral status and moral rights and claims. However, I will need to assume that

**Assumption II** Two beings of the same moral status may enjoy different rights and claims.

This assumption is shared by those mentioned above who have discussed the objection that I address in this article. It is also consistent with common sense, which allows, for example, that two beings with the same moral status may have different rights or claims because of their different internal or external circumstances. For example, a more severely injured person may have a stronger claim to medical assistance than a less severely injured one even though they share the same moral status. Or one mildly injured person may have a stronger claim to medical

assistance than a moral equal because her medical assistance can be provided at lower cost.

In addition to these assumptions about moral status, I also need to make some assumptions about the manner and circumstances in which beings of supra-personal moral status would be created, if they were to be created at all. I assume that supra-persons would be created through the enhancement of pre-existing human persons, not *de novo* (for example, by constructing sophisticated computers) nor via the enhancement of other beings (such as non-human animals or human embryos, if these do not qualify as persons). I also assume that these enhancements would not be heritable (the offspring of supra-persons would, in the absence of further enhancements, be mere persons), and that they would be undergone voluntarily.

### 3 The possibility of creating supra-persons

With these assumptions in hand, we can now turn to the objection under consideration. This begins from the claims that

- (1) The technological enhancement of human mental capacities could result in the creation of supra-persons

and

- (2) The creation of supra-persons would harm ordinary, unenhanced humans.

In the next two sections I will assess (1) and (2) in the light of Buchanan's critique of them. In the subsequent sections, I offer and defend my own response to the objection.


First, then, claim (1). Buchanan doubts though does not outright deny (1); he doubts that beings created through the enhancement of humans could have supra-personal moral status. The view on which he leans most heavily in raising these doubts is the view that

*(Threshold)* There is a threshold of mental capacity above which all beings have the same moral status and above which currently typical adult humans lie (Buchanan 2009, pp. 247, 357, 366–367).<sup>6</sup>


This view, which Buchanan takes to be familiar from the Kantian ethical tradition, is consistent with each of the possibilities schematically depicted in Fig. 1. We are to imagine that moral status is depicted on the vertical axis, while mental capacity, or some particular aspect of it, is depicted on the horizontal axis.<sup>7</sup>

<sup>6</sup> See also Savulescu (2009, pp. 237–238), Wikler (2009, p. 346), McMahan (2009a, pp. 601–602), Wilson (2007), Williams (2008, p. 148). I use the term 'mental capacity' rather than Buchanan's favoured 'cognitive capacity' in recognition of the fact that some regard noncognitive mental capacities as determinants of moral status.

<sup>7</sup> Depicting moral status and mental capacity as continuous variables in this way is not without problems, but I hope that the graphs convey the general idea of *Threshold*.



**Fig. 1** Some possible relations between mental capacity and moral status according to *Threshold*



**Fig. 2** Some possible relations between mental capacity and moral status according to *No Threshold*

In the first graph, there are a number of thresholds of capacity such that, as one crosses each threshold, moral status rises discontinuously, whereas elsewhere it is constant. In the second graph, moral status rises gradually until it reaches some threshold of capacity beyond which it remains constant. Finally, in the third graph, moral status is a binary variable.<sup>8</sup> In each case, there is some level of capacity below that possessed by typical adult humans, and beyond which moral status no longer rises.

Buchanan contrasts *Threshold* with another view, according to which

(*No Threshold*) Moral status is a strictly and continuously rising function of mental capacity, or some particular mental capacity (such as the capacity for wellbeing).

Each of the possibilities depicted in (Fig. 2) is consistent with *No Threshold*.

If *Threshold* is correct, then enhanced beings could not have greater moral status than ordinary humans: there simply exists no higher moral stratum. *Threshold* thus provides a sturdy bulwark against (1). On the other hand, *No Threshold* does not rule out the creation of supra-persons; it is consistent with (1).

<sup>8</sup> For discussion of these and other views about how moral status varies across different types of beings with lower moral status than persons, see DeGrazia (2008) and McMahan (2008).

Buchanan argues that we should prefer *Threshold* to *No Threshold* and tentatively infers that human enhancement could not create supra-persons. He cites two main advantages of *Threshold*. First, it is easily able to accommodate the widely held intuition that all existing humans who qualify as persons have the same moral status, whatever their mental capacity (2009, pp. 347–348, 359). For example, the most intelligent humans have no greater moral status than those whose cognitive capacities are relatively meager. *Threshold* can account for this moral equality among human persons; it allows us to accept “equality of moral status of the sort we associate with personhood” despite the fact that, among persons, there are “many inequalities, including inequalities in the very capacities that confer moral status” (2009, p. 359). It is doubtful whether *No Threshold* can do the same.

*Threshold* can also accommodate another widely held intuition: that the moral status enjoyed by persons is special (2009, pp. 360–362). According to *Threshold*, the moral status enjoyed by persons may be substantially greater than that enjoyed by beings with lesser mental capacity; there may be a sharp drop-off in moral status below the threshold. On *No Threshold*, on the other hand, persons simply occupy a point or range on a continuum.

However, even if *Threshold* is superior to *No Threshold*, it will not follow that that human enhancement could not create supra-persons. Buchanan claims that unless one adopts *No Threshold* “the worry that biomedical enhancements for some but not all would create a new moral status of [supra-person] is highly dubious”, or, as he later puts it, “very implausible” (2009, pp. 368–369). His thought appears to be that rejecting *No Threshold* forces us into the arms of *Threshold*. But there are other possible relations between mental capacity and moral status, including some that are both plausible and consistent with the possibility of creating supra-persons. Consider, for example, the view that

(*Plateau*) There is a range of mental capacity within which all beings have the same moral status, and within which all currently typical adult humans lie, but above which moral status rises, either gradually or in steps.<sup>9</sup>


This view is consistent with each of the possibilities depicted in Fig. 3.

*Plateau* can plainly accommodate the moral equality of all human persons: we may all lie on the ‘plateau’. It can also accommodate the view that the moral status enjoyed by persons is special in the sense that it is substantially greater than that enjoyed by less mentally able beings. But it allows that enhanced beings could have greater moral status than persons. Unlike persons, these enhanced beings might lie beyond the right-hand edge of the plateau.

It might be objected that *Plateau* lacks the theoretical underpinnings that can be adduced in support of *Threshold*. *Threshold* is consistent with what Buchanan calls respect-based accounts of moral status, those familiar from the Kantian ethical tradition.<sup>10</sup> These accounts, as understood by Buchanan, have three important features. First, they hold that there is some mental capacity that confers a special

<sup>9</sup> A similar view has also been entertained by McMahan (2009a, pp. 601–602).

<sup>10</sup> See, for example, Buchanan (2009, pp. 357, 360–361), Rawls (1971, esp. pp. 504–512), Korsgaard (1986, esp. p. 93), and Darwall (2006). For discussion, see, Arneson (1999, pp. 118–122).



**Fig. 3** Two possible relations between mental capacity and moral status according to *Plateau* (The dotted line to the left of the lower threshold denotes that, like *Threshold*, *Plateau* is silent on the relation between mental capacity and moral status in this region.)

moral status, and that is possessed by currently typical adult humans. This may be, for example, the capacity for practical reasoning, for moral agency, or for engaging in practices of mutual accountability. Second, they hold that the *degree* to which one possesses that capacity is unimportant to moral status; it matters only that one possesses the relevant capacity to some non-zero degree (which is what I will henceforth mean when I say that a being possesses a capacity without further specifying its degree of capacity). Third, as Buchanan interprets them, respect-based accounts include the assumption that there is no higher capacity, not possessed by currently typical adult humans, that could confer higher moral status. It follows from these three features that typical adult humans have *maximal* moral status, as do any other beings which possess the relevant capacity. Possession of that capacity constitutes a threshold beyond which moral status no longer rises.

Respect-based accounts of moral status thus support *Threshold*. I am aware of no popular competing account of moral status that supports *Plateau*. However, it is easy to imagine a variant of the respect-based account that would support it. Consider an account according to which two mental capacities, *F* and *G*, are all that matter for moral status. A being with both of these capacities has full moral status. A being with only one has lesser moral status, and a being with neither has none. There are thus three discrete tiers of moral status. Suppose further that Capacity *F* is possessed by no non-human animals but by all currently typical adult humans. Perhaps *F* is the capacity for engaging in practices of mutual accountability, for practical reasoning, or for moral agency. Capacity *G* is possessed by no currently existing humans or animals: it is a higher capacity—perhaps the capacity for higher forms of co-operation or normative deliberation or moral agency. A being that possessed both *F* and *G* would, on this view, have greater moral status than ordinary humans, who possess only *F*.

As far as I am aware, no philosopher has systematically defended a three (or more) tiered account of moral status in which humans occupy any other than the highest tier.<sup>11</sup> But this has little evidential significance. After all, philosophical work

<sup>11</sup> Though Jeff McMahan has expressed openness to such a view by defending multi-tiered views of moral status while acknowledging the possibility of a tier higher than that occupied by persons. See especially McMahan (2008, pp. 93–104).



on moral status has typically been motivated by a desire to explain the elevated or equal moral status of humans, and, in some cases, to accommodate the common sense view that some animals also have significant moral status. If these are one's goals, there is no need to consider, indeed no point in considering, higher capacities that no animals or humans actually possess. The absence of such accounts from the literature thus provides little or no evidence for their falsity. There is, moreover, reason to consider such multi-tiered accounts as serious alternatives to respect-based accounts. There is reason to be suspicious of the presupposition, made by respect-based accounts, that capacities lacked by typical adult humans are irrelevant to moral status; it would be a surprising good fortune for humanity if the threshold for maximal moral status lay just below the level of mental capacity typical of ordinary adult humans.

Buchanan might reply, at this point, that existing respect-based accounts have not, as I suggest above, been developed in order to explain the moral status of humans and non-human animals. Rather, these accounts have fallen out of systematic moral theories, of a broadly Kantian nature, that are plausible (partly) for independent reasons. I cannot engage in an appraisal of Kantian moral theory here. I hope it suffices to note, first, that there is reasonable disagreement about whether Kantian theories are indeed plausible, and second, that multi-tiered accounts of moral status of the sort that would support *Plateau* might in fact be quite consistent with the most attractive elements in those theories. For example, these multi-tiered accounts are not obviously inconsistent with the ideas that morality requires acting on universalizable maxims, that rational agents ought to be treated as ends in themselves, and that the moral law is created by rational agents.

Alternatively, Buchanan might reply that it is difficult to call to mind any capacity that could plausibly confer supra-personal moral status. It is difficult to think of a mental capacity, lacked by ordinary humans but potentially possessed by enhanced beings, that could fill the role of Capacity *G* in the three tiered account that I sketched above. According to respect-based accounts of moral status, persons have greater moral status than lower beings not because they possess the same capacities to a greater degree, but because they possess entirely new capacities. They are not just better at performing the mental tasks that lower beings perform; they can do entirely different things, such as engaging in moral reasoning. Perhaps enhanced beings would, likewise, need to possess a new and qualitatively different mental capacity in order to enjoy supra-personal moral status. As Buchanan notes, it is not clear what that capacity could be (2009, pp. 359, 362–363).

However, we should not infer from the fact that it is difficult to call to mind such a capacity that none could exist. We may, as Buchanan admits, simply be faced with a “failure of imagination” (2009, p. 359). If not for the fact that some people already possess them, it would be difficult to call to mind many capacities that enhanced beings clearly could possess. I doubt that I would speculate that enhanced beings might possess a capacity to perform mental manipulations on objects in more than four dimensions, or to correctly identify lone musical notes, were I not aware that some people already possess this ability.

In any case, it may be easier than it seems at first sight to imagine capacities that would plausibly confer supra-personal moral status. Jeff McMahan has argued that

substantial enhancements of a number of different mental capacities might result in the emergence of qualitatively new capacities in much the same way that typically human capacities such as reasoning and sophisticated language abilities emerged as *Homo sapiens* evolved from great apes, which lacked those capacities. He speculatively considers, as an example of such an emergent capacity, that enhanced beings might acquire novel means of experiencing the mental states of others:

It is widely held that empathy is relevant to, and perhaps even necessary for, moral agency, and many philosophers have held that the capacity for moral agency is necessary for the higher form of moral status. Suppose, then, that supra-persons would have a capacity that would be better for moral agency than mere empathy. Suppose they could actually experience other individuals' mental states while simultaneously reflecting on those experiences in a self-conscious manner from their own point of view. This would require a divided form of consciousness, but that would be only a rather extreme instance of the fragmentation of consciousness of which we are increasingly aware in ourselves (McMahan 2009a, p. 604).<sup>12</sup>

It is perhaps plausible to think that beings with this new capacity would possess a higher form of moral agency than ordinary humans, and that this would endow them with a higher level of moral status.

I have a further suggestion, and one that may tax our imaginative abilities less than McMahan's. Perhaps supra-personal moral status would be conferred by the capacity for constructive participation in some new form of social co-operation. It is not inconceivable that quantitative increases in capacities for altruism, self-control and general intelligence might lead enhanced beings to develop new and qualitatively different forms of social co-operation. For example, these beings might replace existing political and legal institutions with an anarchic system in which individuals are assumed sufficiently altruistic and intelligent to avoid behaviors that contribute to collective action problems. Or they might replace existing economic markets with ones in which prices are determined by an assessment of a fair price by the buyer. Plausibly, we unenhanced humans are incapable of constructive participation in such forms of co-operation, but enhanced beings might differ from us in this regard. Quantitative increases in certain mental (including moral) capacities might confer the entirely new capacity for constructive participation in such co-operative activities. Moreover, it is conceivable that this new capacity would confer supra-personal moral status. After all, the *inability* of

<sup>12</sup> A recent case of conjoined twins adds credibility to McMahan's suggestion. Krista and Tatiana Hogan, born in 2006, are joined at their heads and their thalami are thought to be connected by a neural pathway. One of the functions of the thalamus is to relay and process sensory information, and the connection in these twins appears to allow one to recognize some of the sensory inputs of the other. In one experiment, a strobe light flashed in Tatiana's eyes elicited a marked electrical response in Krista's occipital cortex, the brain area responsible for processing visual inputs. The parents also report that one twin will laugh at television images which only the other is viewing. On occasion, the twins also appear capable of responding to the non-sensory mental states of one another, such as thirst, without verbal communication. If these twins can genuinely experience each other's mental states in virtue of their neural connection, one might speculate that it will, in the future, be possible to engineer similar capacities by creating an artificial pathway. For a recent discussion of this case, see Dominus (2011).

non-human animals to constructively participate in characteristically human forms of co-operation (democracy, economic markets, the operation of legal systems and so on) is one of the more plausible (or less implausible) grounds for assigning those animals sub-personal moral status.

Of course, even if neither this suggestion nor McMahan's is plausible, it remains possible that there is some mental capacity that we cannot, due to our own lesser capacity, currently imagine, but that would confer supra-personal moral status. Our inability to call such a capacity to mind would not rule out the possibility that one could exist.

#### 4 The harm of creating supra-persons

In addition to questioning the suggestion that human enhancement might produce supra-persons, Buchanan also questions (2), the claim that the creation of supra-persons would harm ordinary, unenhanced humans. In assessing this claim I will assume, as Buchanan does, that some mere persons would still exist following the creation of supra-persons (not *all* would make the transition to supra-personhood), and that the supra-persons would live among and interact with these mere persons.

An obvious reason for denying (2) is that moral status is not zero-sum: one being's gain in moral status does not necessitate a loss for any other. So even if enhanced humans acquired supra-personal moral status, ordinary humans would retain their existing moral status—that associated with personhood. However, as Buchanan notes, even if one's moral status would not be diminished by the presence of beings with greater moral status, the *value* of that moral status for the individual might be diminished (2009, pp. 349, 363–364). One explanation for why moral status might be devalued in this way is that it is a partly positional good: its value depends in part on one's relative, rather than absolute, endowment of it. Consider the following scenario.<sup>13</sup> Suppose, as is somewhat plausible, that it is morally permissible for a person to kill one non-human primate if and only if this will (a) prevent the death of *three* or more primates of *equal* moral status, or (b) prevent the death of *one* or more primates of *greater* moral status. Now imagine a population of ten primates all with the same moral status. Each can be permissibly sacrificed (that is, intentionally killed) to prevent the death of three or more others, but none can be permissibly sacrificed to avert the death of one or two others. But now suppose five of the primates acquire a higher moral status through cognitive enhancement. Although the unenhanced primates have the same moral status as before, they are now more susceptible to permissible sacrifice. Whereas before, they could be sacrificed only to prevent the death of three or more others, they may now, pursuant to (b), be sacrificed merely to avert the death of one or two of their enhanced contemporaries. The moral status of the unenhanced primates has diminished in value. Previously, in addition to any other value that it may have had, their moral status had the instrumental value of conferring a high degree of

<sup>13</sup> The example is modified from Buchanan (2009, pp. 363–364).

immunity to permissible sacrifice. Now it confers less immunity.<sup>14</sup> The unenhanced primates have thus, plausibly, been harmed.

One could, perhaps, deny that the loss of immunity to permissible sacrifice suffered by the unenhanced primates in this case constitutes a harm. After all, the loss may make no difference to how their lives go. Arguably, what matters is actually avoiding sacrifice (permissible or otherwise), not being immune to it. Still, being rendered less immune to permissible sacrifice can reasonably be expected to increase the likelihood that one will indeed be sacrificed. Thus, it constitutes at least an *expected* harm, and expected harms might be thought to qualify as harms themselves, or at least to be morally equivalent to harms.

In any case, Buchanan appears to accept that non-human animals might be harmed through losing immunity to permissible sacrifice (2009, p. 364). What he resists is the suggestion that *persons*—such as ordinary, unenhanced humans—could be harmed in this way. This resistance is based on the thought that being a person makes one *inviolable*, and that our inviolability protects us against permissible sacrifice of the sort that befalls the unenhanced primates.

McMahan (2009a, pp. 601–602) has noted a difficulty with this response to (2): our inviolability, as persons, does not protect us against *all* permissible sacrifice. To say that persons are inviolable is to say that they have certain rights—for example, a right not to be intentionally killed, a right to determine what is done to one's body and so on. But these rights are not, most would agree, absolute and unexceptionable. They may be overridden in, or may fail to extend to, extreme situations. Even if I am inviolable and therefore have a right not to be intentionally killed, it may be permissible to intentionally kill me in extreme circumstances, say, when the only alternative involves the death of many more people. It might also be permissible to intentionally kill me when the only alternative involves the death of a being *more inviolable* than me. Thus, even if persons are inviolable, they might be rendered more susceptible to permissible sacrifice by the creation of beings with *greater* inviolability.

Buchanan responds to this worry by claiming that

... inviolability, properly understood, is a threshold concept... meeting the requirements for being a person confers inviolability and that is what counts; having the characteristics that confer personhood to a higher degree does not confer greater inviolability. Hence, whatever exceptions there are to the assertion that persons may not be sacrificed for the sake of other persons apply equally to all who qualify as persons, mere persons and [supra-persons] included.<sup>15</sup>

The thought here is that persons are not merely inviolable, they are *as inviolable* as supra-persons. Thus, we could not be permissibly sacrificed for the sake of supra-persons by reason of our lower inviolability.

<sup>14</sup> In saying that the unenhanced primates now have less immunity to permissible sacrifice, I do not mean that they may now be sacrificed for the sake of less important goals; I mean that they may be sacrificed in a wider range of circumstances. The reason that they can be sacrificed in a wider range of circumstances is precisely that some goals for which they previously could not be permissibly sacrificed—the saving of certain other primates—have increased in importance.

<sup>15</sup> Buchanan (2009, p. 364). For further elucidation and defence, see pp. 364–367. Though Buchanan subscribes to this claim, he does not claim to have conclusively proven it. See esp. p. 366.

There is a further problem, however. There are cases in which one being might be permissibly sacrificed for the sake of another even though both are inviolable, and equally so. Consider a case in which you face a choice between killing *A* and killing *B*, where *A* and *B* are fully and equally inviolable but *A* has a much stronger interest in survival, say, because *B* is about to die anyway. (Suppose that if you do not kill one of them, both, along with many others, will be killed by someone else.) Plausibly, it would be permissible to kill *B* but not *A* in this scenario. The reason: considerations of inviolability are tied, and in such circumstances, interest in survival serves as a tie-breaker. Though our maximal inviolability protects us against being permissibly sacrificed for others by reason of our lesser inviolability, it does not protect us against being permissibly sacrificed when we have weaker claims to continued survival for other reasons—for example, because we have weaker interests in continued survival—and considerations of inviolability are tied. This is significant because it seems possible that supra-persons might, in general, have stronger claims to continued survival than mere persons. They might have stronger interests in continued survival than us. Or their interests might matter more than ours. Either way, their existence might render mere persons more susceptible to permissible sacrifice in tie-breaker cases.

In addition, there are varieties of permissible harm other than sacrifice (i.e., intentional killing) against which our maximal inviolability provides no protection at all. Even if I, as a person, am inviolable, and maximally so, it would be permissible for a charitable or public hospital to deprive me of headache cure in order to instead cure some other person with a worse or more easily curable headache. This is because my claim to the cure derives not from my inviolability, but from my interest in being free from headaches, and, as it happens, my competitor has a stronger interest of the same sort. Supra-persons, one might think, could have exceptionally strong claims to having even mild headaches cured (perhaps because they have exceptionally strong interests in receiving such cures, perhaps because their interests count more). Thus, if supra-persons existed, it might be permissible to deprive me of a headache cure in order to cure a much milder headache endured by a supra-person. That I am inviolable, or as inviolable as the supra-person, is neither here nor there in this context.

The more general point is that there are kinds of permissible harm not involving intentional killing against which our inviolability confers no protection. It thus does nothing to prevent the creation of supra-persons from reducing what immunity we do have to these harms. And this reduction in immunity to permissible harm might itself count as a harm—a kind of *meta*-harm.

## 5 The morality of creating supra-persons

It is difficult to confidently reject either (1), the claim that the technological enhancement of human mental capacities might create supra-persons, or (2), the claim that the creation of supra-persons would harm ordinary humans. These claims may be false, but we have discovered nothing that clearly demonstrates this.

Suppose that (1) and (2) are in fact true. Unrestrained pursuit of enhancement might then, through the creation of supra-persons, harm ordinary humans. Plausibly,

we have moral reasons not to bring about this harm. Thus, it appears we have moral reasons not to encourage at least some human enhancements: those that would result in the creation of supra-persons (henceforth, ‘status enhancements’). This is precisely what (3) maintains.

It is true that the harm to unenhanced individuals caused by others undergoing status enhancements is quite different from the sort of harm usually considered by moral philosophers. It is a second order harm: the harm of having one’s immunity to other permissible harm reduced. Nevertheless, it seems plausible that we typically have reason not to bring about second order harms, just as we typically have reasons not to bring about first order harms, particularly since being rendered less immune to permissible harm will typically increase the likelihood that one suffers a first order harm.

So I think that (1) and (2) may be correct, and that they provide some support to (3). However, I doubt that this should concern the proponent of human enhancement. In the remainder of this section, I outline why.

First, reasons not to impose the second order harm of reduced immunity to permissible harm (henceforth sometimes just ‘immunity’) are not always decisive. Suppose, as is plausible, that emergency medical staff have stronger moral claims than others to the avoidance of certain harms. Perhaps they have stronger claims to exemption from jury duty and to the waiving of charges for vaccination services. Imagine a society in which there are no emergency medics. A politician then decides to train and employ some. In doing so, that politician reduces the immunity of those who do not become emergency medics to certain permissible harms. If a situation arises in which Arnold or Bridget must do jury duty, it is more likely to be permissible to impose the harm on Arnold now that Bridget has become an emergency medic than it would have been had she not. However, it is not clear that the harms imposed on those who remain non-medics count decisively against the politician’s measure. Perhaps she had some reason not to train emergency medics because of the reduction in immunity that this would cause to others (I am unsure about this), but it is certainly not an invariably decisive reason. It might be outweighed by countervailing reasons—for example, reasons to improve healthcare provision.

The case just described is most naturally understood as one in which the politician reduces the immunity of some to permissible harm not by increasing the moral status of others, but by changing their circumstances. However, similar thoughts may apply to cases in which one reduces the immunity of some by bringing it about that others become supra-persons, for example, by encouraging status enhancements. The consequent reduction in the immunity of the unenhanced humans may generate *some moral reason* not to encourage status enhancements, but, as in the case of training emergency medics, this reason need not be decisive. For example, it might be outweighed by countervailing reasons *to* encourage such enhancements, say, because doing so would increase overall human productivity or accelerate the rate of scientific progress.<sup>16</sup>

---

<sup>16</sup> For discussion of the potential social benefits of human enhancement, see, for example, Buchanan (2008), Bostrom and Roache (2011); Vedder and Klaming (2010).

In fact, we might have good moral reasons to encourage status enhancements even if doing so would have no economic or scientific benefits. This is because, just as the status enhancements would reduce the immunity of, and thereby harm, unenhanced humans, so too they would *increase* the immunity of, and thereby benefit, those who make the transition from personhood to supra-personhood. The loss of immunity endured by the unenhanced humans would be matched by a gain in immunity enjoyed by those who undergo the enhancements. Thus, though we may have some reason not to encourage status enhancements, because of the reduction in the immunity this would impose on the unenhanced individuals, we may also have a precisely parallel reason *to* do so.

What we are faced with here is a possible *redistribution* of immunity. Status enhancements would not result in any net decrease in immunity, at a population level. Immunity would simply be redistributed from those who remain mere persons to those who become supra-persons. The important question is whether we have most moral reason not to encourage status enhancements given the redistribution of immunity that would result.

In what follows, I assess four attempts to demonstrate that we indeed have most reason not to encourage status enhancements.

### 5.1 Inequality of immunity

Status enhancements by some but not others would introduce new *inequalities* in immunity to permissible harm across different beings. In particular, they would introduce inequalities between unenhanced humans and the enhanced individuals. Thus, one way of objecting to such enhancements would be to maintain that there is value in equality of immunity across beings. If this were so, then status enhancements would in one way worsen the overall pattern of distribution of immunity, and it might follow that we have most reason not to encourage these enhancements.

It is doubtful that there is any value in equality of immunity, however. Indeed, one cannot maintain that inequality of immunity exists without that inequality being, in a sense, appropriate or justified. For immunity is a normative notion. One being has greater immunity to permissible harm than another just in case it enjoys stronger moral claims against certain kinds of harm, and that cannot be so if the difference in the strength of those claims is unjustified.

Thus, most believe that there is currently a grossly unequal distribution of immunity across different kinds of beings such that humans typically possess much more immunity than pigs, which possess much more immunity than chickens, which possess much more immunity than scallops. If these differences in immunity indeed exist, this must be because there are further differences, for example in moral status or morally relevant circumstances, that justify them. We may have difficulty spelling out these further differences,<sup>17</sup> but if we are to maintain that there is inequality in immunity, the differences must be there somewhere. If, say, there is no

---

<sup>17</sup> For detailed discussion of some of these difficulties, see Singer (1990), McMahan (1996, 2002, pp. 204–232, 2008), Arneson (1999), and Vallentyne (2005).

morally relevant difference between a human and a chicken, then the human is *not* more immune to permissible harm than the chicken after all.

Similarly, if status enhancements give rise to further inequalities in immunity, those new inequalities must also be, in a sense, justified. They are justified by the very differences in moral status which produce them. Since immunity to permissible harm is a normative notion, one being cannot have greater immunity than another unless the difference is justified. So a new difference in moral status could not give rise to a new inequality in immunity unless it also justified it.

By analogy, suppose that differences in *culpability* between persons give rise to differences in their *moral liability to punishment*. Then if there is a difference in culpability between two persons, there will be an inequality in their liability to punishment. But this inequality arises from a difference in culpability which also justifies it.

Now suppose that two individuals are initially equally culpable, but one then does some wrongful act which makes him more culpable than the other. A new inequality in liability to punishment will have been produced. But, as with existing inequalities, that inequality will be justified, and for the same reason. The difference in culpability which produces it also justifies it. There has been a change in the distribution of liability, and it is a change in the direction of greater inequality. But in a sense this change is not an adverse one, since the inequality in liability maps on to (and is indeed produced by) a difference in culpability that justifies it. Similar thoughts apply, I believe, to inequalities in immunity to permissible harm generated by status enhancements.

Admittedly, there is perhaps a derivative sense in which the inequality in immunity created by some persons undergoing status enhancements could be disvaluable *even though* it would map on to differences in moral status that justify it. It could be disvaluable if the distribution of mental capacities that grounds moral status, and thus immunity, were itself disvaluable.<sup>18</sup> Suppose that the technologies enabling status enhancement are monopolized by a group of scientists who ensure that all members of a certain racial or socio-economic group are prevented from accessing these technologies. As a result, though many others would like to, only members of the privileged group(s) actually undergo status enhancements. In this case, it seems plausible that we wind up with an unjust distribution of mental capacity across different beings. It seems unjust that some possess greater mental

<sup>18</sup> Similarly, it might be argued that existing inequality in mental capacity across animals (including humans) of different species is disvaluable. Vallentyne (2005) tentatively argues that egalitarian concerns could require enhancement of the mental capacities of non-human animals. McMahan (1996) takes the same possibility seriously. The suggestion that existing cross-species inequality in mental capacity is disvaluable would be most plausible if some non-human animals qualify as persons. These animals would then, along with human persons, fall straightforwardly within the sphere of egalitarian justice, which might well deem the unequal distribution of mental capacity across persons to be unjust. It might seem less clear that inequality in mental capacity could be disvaluable if no non-human animals are persons; the sub-personal moral status of non-human animals may seem to disqualify them from the sphere of egalitarian justice. But egalitarian justice may extend to cover beings of somewhat lower moral status than persons. Moreover, the distribution of mental status across beings might be disvaluable for reasons other than that it is unjust. Thus, even if one supposes that all non-human animals are sub-persons, the distribution of mental capacity across existing animals might be disvaluable. So, derivatively, might the resulting distributions of moral status and immunity to permissible harm.



capacities than others as a result of this discriminatory policy. Perhaps this injustice in the distribution of mental capacity across beings would infect the resulting distributions of moral status and immunity.

It is, however, doubtful whether status enhancements by some, but not others, would inevitably lead to an unjust or otherwise disvaluable distribution of mental capacity across beings. Suppose, for example, that access to status enhancements is decided via a lottery procedure to which all agree in advance. Now, it seems doubtful whether the resulting distribution of mental capacity would be unjust or otherwise disvaluable. Or suppose status enhancements were made readily available to all, but that only some chose to use them, with others preferring to remain mere persons. Again, it seems doubtful whether resulting differences in mental capacity would be disvaluable. So it is difficult to see how the resulting distribution of moral status, and thus immunity, could be criticized.<sup>19</sup>

## 5.2 Partiality

I have acknowledged that we may have some reason not to encourage status enhancements: doing so would reduce the immunity of those who remain unenhanced. But if we have that reason, we plausibly also have a parallel reason *to* encourage these enhancements: doing so would increase the immunity of those who undergo the enhancement. There would in fact be no net change in total immunity, only a redistribution of it.

In the previous subsection, I considered the possibility that one might nevertheless argue against status enhancements by maintaining that they would lead to a worse overall pattern of distribution of immunity. I took an impartial point of view, evaluating the distributive pattern of immunity that would exist in the post-enhancement world—that is, the world in which some persons had become supra-persons. Perhaps this was the wrong approach, however, and perhaps deviating from it would open up new possibilities for opposing status enhancements.

An alternative approach would focus not on the pattern of distribution of immunity that would result from status enhancements, but on the losses and gains in immunity that these enhancements would impose or confer on individuals. This approach would open up further means of objecting to status enhancements. Although the loss of immunity endured by those who remain mere persons would in a sense be matched by the gains in immunity accruing to those who become supra-persons, it could be argued that we should care more about the losses than the gains.

One way of defending greater concern for the losses than the gains would appeal to the idea that we should be *partial* towards those who suffer the losses of immunity—we should attach greater weight to the losses of immunity than to the

---

<sup>19</sup> Note also that, if the reason for thinking that greater inequality in immunity would be disvaluable is simply that the inequality of mental capacity on which it is based is disvaluable, then we are no longer dealing with a distinctively moral status-based concern about human enhancement. The basic concern is now simply that enhancement by some but not others would result in disvaluable inequalities in mental capacity. The implications of this inequality for moral status and immunity to permissible harm may exacerbate the concern—we may care more about inequalities in mental capacity if we know that they will also generate inequalities in moral status and immunity—but the problem would be there regardless.

gains of immunity because we should care more about ordinary, unenhanced human persons (henceforth, ‘the unenhanced’) than about the enhanced supra-persons (henceforth, ‘the enhanced’). It is certainly natural to empathize more with the unenhanced than the enhanced: the unenhanced are more like us. Perhaps it could also be argued that we *should* care more about the unenhanced, rather as it is sometimes argued that parents should care more about their own children than others, and that citizens should care more about their compatriots than others.

This approach seems unpromising, however. An initial problem is that, at the time when a status enhancement would take place—which is also the time at which the gains and losses of immunity would occur—*all* are unenhanced. There are, at that point, no enhanced individuals to be partial against.<sup>20</sup> Thus, even if we have reason to be partial towards the enhanced, and against the unenhanced, those reasons do not yet apply at the time that a status enhancement occurs.

It might be responded that in deciding whether to encourage status enhancements, one could nevertheless be partial against the enhanced, though in a non-standard kind of way. One could be partial in the sense that one cared less about gains and losses of immunity when those who received those gains and losses would, *after receiving them*, be enhanced individuals. This would, however, be an unappealing kind of partiality. It would be akin to the partiality that a doctor would exhibit if she cared less about the effect of an intervention on her patient *A* than about the effect of another intervention on her patient *B* on the basis that patient *A* would no longer be her patient following the intervention, whereas *B* would remain her patient. Intuitively, the fact that *A* is a patient of the doctor at the time the intervention is performed is enough to for her to fall within the scope of any partiality that doctors have reason to show towards their own patients.

In any case, even if we accept that this non-standard kind of partiality could be supported by moral reasons, it is difficult to see any reason why it *would* be so supported in the present case. It is difficult to see why the kinds of relations that hold between us and future unenhanced humans could support partiality, of either a standard or non-standard variety, towards them. We clearly have no *relationship* with these individuals of the sort that parents have with their children or doctors have with their patients. Certain significant relations do obtain between us and the unenhanced. For example, we existing humans stand in the relation of species co-membership with future unenhanced humans. However, it is highly controversial whether species co-membership can ground moral reasons (or even moral permissions) to be partial. Moreover, the enhanced beings—the supra-persons—might also be humans: it is not clear that even dramatic enhancements of mental capacity would have to be species-altering. Thus, even if humans should have greater concern for one another than for others, this may not give us grounds to care more about the unenhanced than the enhanced.

It could be argued that merely bearing the relation of moral equality to another being—sharing one’s moral status with it—is enough to give one reason to be partial to that being. On this view, persons should be partial to other persons, and supra-persons to other supra-persons. It is perhaps somewhat more plausible that the

<sup>20</sup> I thank Jeff McMahan for raising and pressing me to address this issue.

relation of moral equality could generate reasons to be partial than it is that sharing one's species membership generates similar reasons. Moral equality is, after all, morally significant in a way that species co-membership is not. However, it remains obscure why the relation of moral equality should be attached similar significance to being in a parent–child or doctor–patient relationship. The relation of sharing moral status involves none of the interaction, mutual expectation or recognition of a special connection that characterize those relationships.

In addition, the view that we should have greater concern for those who share our moral status has some unappealing implications. For example, it implies that in addition to our reasons to treat animals less well than ordinary humans, since they have lower moral status, we have a further reason to treat them less well: they have *different* moral status and therefore fall outside of the scope of our reasons to be partial. This is at least mildly counter-intuitive. Intuitively, if it is indeed permissible to treat animals less well than humans, it is their *lower* moral status, not their different moral status, that does all the normative work.

### 5.3 Harms versus benefits

A more promising way of arguing that we should attach greater weight to the losses of immunity caused by status enhancements than to the corresponding gains would appeal not to partiality, but to the claim that we have, quite generally, stronger reasons not to contribute to harms—normally understood as losses—than to contribute to benefits—normally understood as gains. This view is often said to be a part of common sense morality (McMahan 2009b). It can be formulated, in a general way, as follows: for any agent *A*, amount of harm/benefit *x*, and means of contributing to harm/benefit *y*, *A* has greater reason not to contribute in manner *y* to *x*-amount of harm than to contribute in manner *y* to *x*-amount of benefit. If this view is correct, then we have stronger reasons not to reduce the immunity of the unenhanced by pursuing status enhancements than we have to increase the immunity of the enhanced to the same degree, and in the same way. This suggests that, overall, we will have most reason not to bring about the package of losses and gains in immunity that status enhancements would impose.

A problem, however, is that, according to common sense morality, we have some reason to contribute to benefits, though it may be weaker than the reason not to contribute to comparable harms. Moreover, if a benefit exceeds a harm by a wide margin, the reason to contribute to the benefit will, in many cases, be stronger than the reason not to contribute to the harm. Reasons not to contribute to harm are not often of invariably over-riding importance. This suggests that we could have sufficient reason to encourage status enhancements if they would cause a loss in immunity to the unenhanced, a comparable gain in immunity to those who undergo the enhancements, and *further benefits*. It is plausible that enhancements that increased moral status would indeed confer benefits besides the gain in immunity that accrues to the enhanced. For example, they might cause substantial increases in productivity that benefit almost everyone (Buchanan 2008). So it remains unclear whether concerns about the immunity of the unenhanced would always, or even often, give us decisive reasons against encouraging status enhancements.

Admittedly, there is one sort of case in which reasons not to contribute to harm cannot, or cannot in any normal circumstances, be outweighed by countervailing reasons to contribute to benefit: where contributing to the harm would constitute a rights violation. Thus, if it could be shown that encouraging status enhancements would violate the rights of those who remain unenhanced, then we could confidently conclude that there is most reason not to encourage such enhancements.

Do the unenhanced have a right to their immunity that would be violated by encouraging status enhancements? One reason to doubt that they do is that there are cases in which our reasons against contributing to losses of immunity are not, intuitively, decisive. Recall the case of the politician who decides to train and employ emergency medics, thereby reducing the immunity of others to the permissible imposition of jury duty or vaccination charges. If there were a general right to retain one's immunity and a corresponding duty not to contribute to losses of immunity, the politician would surely have decisive reasons not to do this.

Of course, we may have more limited rights, such as rights not to have our immunity reduced in such-and-such a way, or to such-and-such degree, that would invariably be violated by encouraging status enhancements though they are not violated, for instance, by training emergency medics. I see no good reason for supposing that we have such rights, however. I cannot think of any cases eliciting intuitive responses that favor their existence, nor can I see any good theoretical grounding for them. It thus seems safest to conclude that encouraging status enhancements need violate no right to preservation of immunity. If this is correct, then, though we may have somewhat stronger reasons not to contribute to losses of immunity than to contribute to comparable gains, that reason might be outweighed where the creation of supra-persons would also confer other benefits as, plausibly, it often would.

#### 5.4 Identity change

I have been describing status enhancements as transformations in which fixed individuals make the transition from personhood to supra-personhood. One final attempt to show that we have most reason not to encourage status enhancements would question the accuracy of this description. It would appeal to the thought that any human enhancement capable of creating a supra-person would be identity altering. The supra-person would not be the same individual as the person who existed prior to the enhancement. If this is right, then no individual gains immunity as a result of the status enhancement. Rather, a being with less moral status, and thus immunity, is simply replaced by another being with more moral status and immunity. It may be that the supra-person created by the enhancement enjoys a benefit in a noncomparative sense: arguably his high level of immunity could be good for him though he has received no *gain*. There is also a comparative benefit (or gain) from the 'eye of the universe', since one individual with less immunity is replaced by another with more. But no individual is benefitted in the straightforward sense of gaining immunity. On the other hand, there are individuals who are harmed in the straightforward sense of losing immunity. Those persons who remain unenhanced undergo no change in identity; they remain the same people, though

they lose immunity as a result of the enhancement undergone by others. In addition, the person who undergoes the enhancement would typically be harmed by going out of existence, though this is, we are assuming, a voluntarily accepted harm. The enhancement is similar to—or perhaps just is—suicide for that person.

If status enhancements are identity altering, then the overall package of harms and benefits that a status enhancement produces is as follows. (1) There are harms, in the form of losses of immunity, suffered by those who remain unenhanced. (2) There is typically a harm, through going out of existence, for the person who undergoes the enhancement. On the other hand, there is no straightforward benefit in the form of a gain in immunity to any individual, though there may be (3) a noncomparative benefit, in the form of a high level of immunity, enjoyed by the supra-person brought into existence by the enhancement, and (4) a gain ‘from the eye of the universe’ in the form of the replacement of a less immune individual with a more immune one. There may also, of course, be (5) straightforward economic, social or scientific benefits (or indeed harms) to those who remain unenhanced due, for example, to the greater productivity of the enhanced. It is unclear whether this overall package of harms and benefits would be one that we have most reason to encourage. Benefits (3) and (4) are arguably morally less weighty than harms (1) and (2), since they do not consist in gains to fixed individuals, and it is not clear that general economic, social and scientific benefits—(5)—could make up for this difference.

Fortunately, for the proponent of human enhancement, it is very doubtful whether status enhancements would have to be identity altering. It is true that existing philosophical work on our identity tends to address the question ‘what is required for *personal* identity?’ This question is naturally read as presupposing that we are essentially persons; that we could not cease to be persons while remaining the same individual. Moreover, some have explicitly defended the view that we are essentially persons (Baker 1999, pp. 154–155, 2000, pp. 105–106). If this view is correct, then identity could not be preserved across the transition from personhood to supra-personhood. However, when philosophers hold that we are essentially persons, what they presumably have in mind is that our identity could not be preserved across the transition to *sub*-personhood; we could not remain the same individual while losing that which makes us persons rather than sub-persons. It seems doubtful whether those who hold this view would deny that identity could be preserved across the transition from personhood to *supra*-personhood. Indeed, though what follows is necessarily cursory, the dominant accounts of our identity seem quite consistent with the view that a transition to supra-personhood could leave our identities intact. On psychological accounts, identity is preserved when there is psychological continuity, or there are overlapping periods of psychological continuity, across time. Typically, the types of psychological continuity thought relevant are certain kinds of cross-referencing of mental states to one another across time, such as when a mental state at  $t_1$  figures in a memory at  $t_2$ , or where a desire, at  $t_1$ , to have such-and-such mental state at  $t_2$  is indeed satisfied at  $t_2$ . There is little reason to suppose that cross-referencing of these sorts would be threatened by even radical enhancements of mental capacity. The preservation of identity across the transition from personhood to supra-personhood is even more assured on the

dominant alternative account of our identity—the biological account. On this account, our identity is preserved when we survive as organisms. If an enhancement that would supposedly render us supra-persons were to operate via, say, a pharmaceutical or genetic intervention, it would presumably *rely on* our continued survival as organisms.

It seems plausible, then, that identity could be preserved across the transition from personhood to supra-personhood. Of course, it might well be argued that identity could be preserved, but not in such a way as to retain its ordinary significance. Derek Parfit has argued that identity can be lost while what is normatively significant in survival is preserved (1984, pp. 199–280). Arguably, what is normatively significant in survival may also be lost while identity is preserved (McMahan 2002, p. 51). So it is worth noting that, plausibly, everything normatively significant in our survival would also be untouched by the transition to supra-personhood. Again, this can be seen by considering the implications of the dominant philosophical views. Accounts of what matters in survival tend to cite psychological relations similar to those cited by psychological accounts of personal identity: the cross-referencing of mental states. Status enhancements need not threaten those relations. Some authors do suggest more stringent criteria for what matters in survival. For example, some allow that a degree of persistence of one's beliefs, desires or mental dispositions across time may be part of what matters in survival. Clearly, these mental states and dispositions might be altered by radical enhancements of mental capacity (McMahan 2002, pp. 69–81, 316–322). Others hold that for what matters in survival to be preserved across a psychological change, one's life narrative must be preserved; any psychological changes must be folded into a single first-personal life story (Taylor 1989; Schechtman 1996; MacIntyre 1997; DeGrazia 2005). On this view, too, there is the potential for the transition to supra-personhood to disrupt what matters in survival, say, if the necessary enhancements were forced upon people, or were driven by inauthentic desires that were foreign to their narratives. However even on these views, it seems possible that what matters in survival could survive the transition to supra-personhood. We can imagine someone pursuing a series of enhancements in a stepwise fashion based on careful deliberation and in pursuit of a coherent project. She gradually acquires greater and greater mental capacity, and eventually a supra-person emerges. In such a case, there seems little reason to say that a life narrative has been disrupted, and the changes wrought to her mental states and dispositions seem similar to, though ultimately greater than, those brought about through intensive and prolonged education, which is not normally thought to attenuate what matters in survival.

Whether we adopt a psychological or biological account of our identity, it seems plausible that our identity could persist across the transition to supra-personhood. More importantly, it also seems plausible that what matters in survival could survive the transition. If both of these thoughts are correct, then there is no harm through going out of existence to the individual who undergoes the enhancement. Indeed, that individual benefits through acquiring greater immunity in the same straightforward way that unenhanced individuals are harmed through losing immunity. There is no justification for downgrading the benefits relative to the harms by appealing to considerations of identity or what matters in survival.

Of course, we might still downgrade the benefits because we take benefits to be, quite generally, less morally weighty than harms. But as we have seen, it is unclear whether this difference in weight is sufficient to establish that we have most reason not to encourage status enhancements.

## 6 Conclusion

One objection to the technological enhancement of humans mental capacities holds that

- (1) It may result in the creation of supra-persons.
- (2) The creation of supra-persons would harm ordinary, unenhanced humans.
- (3) We have moral reasons not to encourage certain human enhancements: those that would create supra-persons.

Allen Buchanan responds to this objection by questioning (1) and (2), but I have argued that these claims are difficult to undermine. I have argued, *contra* Buchanan, that (1) and (2) could be true even if we reject the view that moral status constantly rises with mental capacity, and even if we allow that persons are inviolable. Moreover, I conceded that (1) and (2) do provide some support to (3).

Nevertheless, the proponent of human enhancement may be able to respond adequately to this objection. The acquisition of supra-personal moral status by some could harm unenhanced humans by reducing their immunity to permissible harm, and this may give us some reason not to encourage status enhancements. But it is not clear that this reason is decisive and I have argued that four attempts to show that it is fail. First, though status enhancements would redistribute immunity to permissible harm in the direction of greater inequality in immunity, this inequality need not be disvaluable. Second, it is difficult to make sense of the idea that we might, in our stance towards status enhancements, be partial to the unenhanced, and against the enhanced. Moreover, even if we could make sense of this idea, it is difficult to see why we *ought* to be partial. Third, though harms may be morally weightier than benefits, so that the losses in immunity caused by status enhancement matter more than the corresponding gains in immunity, it is not clear that this difference in moral importance would be sufficient to outweigh other benefits that human enhancements capable of elevating moral status might have. And fourth, it is plausible that status enhancements need alter neither identity nor what is normatively significant in survival, and if they would not, the gains in immunity that they produce count as benefits to individuals in the same straightforward way that the losses count as harms.

I have assumed throughout that supra-persons would be created through the enhancement of pre-existing persons. If they came about in some other way, for example, through the enhancement of early embryos that do not yet qualify as persons, then some of my arguments would not hold. In that case, there would be no persons who would benefit from becoming supra-persons, though there are others who would be harmed. This would leave room for one to oppose status enhancements on the basis that resulting changes in immunity would harm some

persons without benefitting others. Whether this argument succeeds would depend on the relative importance of person-affecting and other moral considerations, on how bad it is for a person to lose immunity to permissible harm, and on the scale of any scientific, social or economic benefits of status enhancements to those who remain unenhanced.

**Acknowledgments** I thank, for their comments on earlier versions of this article, Jeff McMahan, Allen Buchanan, Katrien Devolder, Ingmar Persson, Dominic Wilkinson and audiences at the Universities of Oxford, Hull, Waikato and Roskilde. These allowed me to correct a number of defects in the paper. I am responsible for the many that remain. I thank the Wellcome Trust (grant number WT087211) for their funding.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution Non-commercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

## References

- Arneson, R. J. (1999). What, if anything, renders all humans morally equal? In D. Jamieson (Ed.), *Singer and his critics* (pp. 103–128). Oxford: Blackwell.
- Baker, L. R. (1999). What Am I? *Philosophy and Phenomenological Research*, 59(1), 151–159.
- Baker, L. R. (2000). *Persons and bodies: A constitution view*. Cambridge: Cambridge University Press.
- Bostrom, N., & Roache, R. (2011, forthcoming). Smart policy: Cognitive enhancement and the public interest. In J. Savulescu, R. ter Meulen, & G. Kahane (Eds.), *Enhancing human capabilities*. Oxford: Wiley-Blackwell. <http://www.nickbostrom.com/papers/smart-policy.pdf>.
- Buchanan, A. (2008). Enhancement and the ethics of development. *Kennedy Institute of Ethics Journal*, 18(1), 1–34.
- Buchanan, A. (2009). Moral status and human enhancement. *Philosophy & Public Affairs*, 37(4), 346–381.
- Buchanan, A. (2011). *Beyond humanity? The ethics of biomedical enhancement*. Oxford: Oxford University Press.
- Darwall, S. (2006). *The second person standpoint: Morality, respect, and accountability*. Cambridge, MA: Harvard University Press.
- de Jongh, R., Bolt, I., Schermer, M., & Olivier, B. (2008). Botox for the brain: Enhancement of cognition, mood and pro-social behavior and blunting of unwanted memories. *Neuroscience and Biobehavioral Reviews*, 32(4), 760–776.
- DeGrazia, D. (2005). *Human identity and bioethics*. Cambridge: Cambridge University Press.
- DeGrazia, D. (2008). Moral status as a matter of degree? *Southern Journal of Philosophy*, 46, 181–198.
- Dominus, S. (2011, May 25). Could conjoined twins share a mind? *New York Times Magazine*. <http://www.nytimes.com/2011/05/29/magazine/could-conjoined-twins-share-a-mind.html>.
- Fukuyama, F. (2002). *Our posthuman future*. New York: Farrar, Straus and Giroux.
- Fukuyama, F. (2004). Transhumanism. *Foreign Policy*, 144(September), 42–43.
- Jenyns, S. (1758). *A free inquiry into the nature and origin of evil. In six letters to –* (4th ed., amended). Dublin: G. and A. Ewing.
- Johnson, S. (1825). Review of Soame Jenyns: A free enquiry into the nature and origin of evil. In F. P. Walesby (Ed.), *The works of Samuel Johnson*. Oxford: Talboys & Wheeler. <http://andromeda.rutgers.edu/~jlynch/Texts/jenyns.html>.
- Korsgaard, C. M. (1986). *Sources of normativity*. Cambridge: Cambridge University Press.
- MacIntyre, A. (1997). The virtues, the unity of a human life and the concept of a tradition. In L. P. Hinchman & S. K. Hinchman (Eds.), *Memory, identity, community: The idea of narrative in the human sciences* (pp. 241–263). Albany, NY: State University of New York Press.
- McMahan, J. (1996). Cognitive disability, misfortune, and justice. *Philosophy & Public Affairs*, 25(1), 3–35.



- McMahan, J. (2002). *The ethics of killing: Problems at the margins of life*. New York: Oxford University Press.
- McMahan, J. (2008). Challenges to human equality. *Journal of Ethics*, 12(1), 81–104.
- McMahan, J. (2009a). Cognitive disability and cognitive enhancement. *Metaphilosophy*, 40, 582–605.
- McMahan, J. (2009b). Asymmetries in the morality of causing people to exist. In M. A. Roberts & D. Wasserman (Eds.), *Harming future persons: Ethics, genetics and the nonidentity problem*. Dordrecht, The Netherlands: Springer.
- Parens, E. (1995). The goodness of fragility: On the prospect of genetic technologies aimed at the enhancement of human capacities. *Kennedy Institute of Ethics Journal*, 5, 141–153.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Clarendon Press.
- Rawls, J. (1971). *A theory of justice* (original edition). Cambridge, MA: Harvard University Press.
- Savulescu, J. (2009). The human prejudice and the moral status of enhanced beings: What do we owe the gods? In J. Savulescu & N. Bostrom (Eds.), *Human enhancement* (pp. 211–247). Oxford: Oxford University Press.
- Schechtman, M. (1996). *The constitution of selves*. Ithaca, NY: Cornell University Press.
- Singer, P. (1990). *Animal liberation* (2nd ed.). New York: Random House.
- Taylor, C. (1989). *Sources of the self: The making of modern identity*. Cambridge, MA: Harvard University Press.
- Vallentyne, P. (2005). Of mice and men: Equality and animals. *Journal of Ethics*, 9(3/4), 403–433.
- Vedder, A., & Klaming, L. (2010). Human enhancement for the common good: Using neurotechnologies to improve eyewitness memory. *American Journal of Bioethics—Neuroscience*, 1(3), 22–23.
- Wikler, D. I. (2009). Paternalism in the age of cognitive enhancement: Do civil liberties presuppose roughly equal mental ability? In J. Savulescu & N. Bostrom (Eds.), *Human enhancement* (pp. 341–355). Oxford: Oxford University Press.
- Williams, B. (2008). The human prejudice. In A. W. Moore (Ed.), *Philosophy as a humanistic discipline* (pp. 135–152). Princeton, NJ: Princeton University Press.
- Wilson, J. (2007). Transhumanism and moral equality. *Bioethics*, 21(8), 419–425.