

Persistent Vegetative State, Akinetic Mutism and Consciousness.

Will Davies and Neil Levy

There is a strong intuition that conscious beings have a special kind and perhaps an especially high degree of moral significance denied to beings that lack consciousness. Though we think that the way in which this special significance is spelt out is often misleading, we think that the intuition is well grounded: some kind of consciousness *does* make a significant difference to a being's moral status.

It is this intuition that explains the excitement surrounding recent work by clinicians and neuroscientists that apparently demonstrates consciousness in patients previously diagnosed as in a persistent vegetative state (PVS).¹ We think that the inference from the experimental data to the presence of consciousness is too hasty.² Building on recent work by Colin Klein, we will suggest that a key assumption underlying this inference is far less solid than it has appeared. Our argument casts doubt on methodologies that take *command following* as a marker for endogenous intentional

¹ It is unclear to us how the putative cases of consciousness should alter our classification of these patients. It might seem that such consciousness would shift these patients from a classification of PVS to the minimally conscious state (MCS). Another possibility, however, is that the degree of consciousness would be sufficient to place these patients in the category of locked-in syndrome. We remain neutral on these classificatory issues.

² One of us has previously questioned whether the data shows that the patients have the kind of consciousness routinely assumed to be at issue, namely phenomenal consciousness. Since phenomenal consciousness is not the kind of consciousness that actually does the bulk of the work in underwriting moral status of the kind and degree at issue, however, that quibble now seems irrelevant.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

agency, and hence for consciousness. Klein himself nonetheless accepts that these patients are conscious, although for reasons other than their apparent ability to follow basic commands. We are not convinced by Klein's reasons for attributing consciousness, and hence adopt a more pessimistic position regarding the interpretation of the data. In closing, we shall argue that even if Klein is right and these patients do enjoy *some* kind of consciousness, this is not a kind of consciousness that could ground the special kind of moral status typically assigned to conscious human beings.

Methodological Assumptions in the Science of Consciousness

Our aim in this section is to briefly review the key findings that have motivated much of the discussion about borderline states of consciousness, and then highlight some key assumptions at work in the relevant scientific studies. We shall not rehearse the empirical evidence in any detail, given that other chapters in this volume provide extensive discussion. Briefly, there are two key pieces of evidence which have impressed commentators. Owen et al. (2006) showed that the neural activation in a PVS patient who was asked to imagine playing tennis and navigating around their own house was very similar to the neural activation exhibited by healthy controls asked to perform the same tasks. Building on these results, Monti et al. (2010) used the same paradigm to develop what was, in effect, an fMRI-based communication system, in which an apparently vegetative patient was able to answer 'yes' or 'no' to questions by imagining playing tennis or imagining navigating a familiar environment.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

The precise details of the neural activation exhibited by PVS patients do not much matter for our purposes. What most interests us are the grounds for the inference from these experimental data to the conclusion that the patients are conscious. The key assumption underlying this inference, with which the experimenters and subsequent commentators largely seem to agree, is that the data provide evidence for the presence of *agency*, which in turn indicates the presence of consciousness (Bayne 2013). More precisely, the patient's capacity to follow task instructions, such as to imagine playing tennis, is assumed to signal an endogenously generated intention to comply with the instructions, which in turn signals the presence of intentional agency, and hence consciousness. We can break down this inference into two components:

Command Following: Command following is a marker of intentional agency; i.e. evidence for command following provides strong evidence for the presence of intentional agency.

Agency: Intentional agency is a marker of consciousness; i.e. evidence for intentional agency provides strong evidence for the presence of consciousness.

In combination, Command Following and Agency seem to guide much of the research into borderline states of consciousness.³ It has been said that first-personal subjective report is the gold standard for the ascription of consciousness. In the absence of such reports, Agency surely provides the *silver* standard for consciousness. Since

³ Indeed one of the core diagnostic criteria for transitioning into MCS from PVS is 'following simple commands' on a 'reproducible or sustained basis' (Giacino et al. [2002] p. 351).

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

command following is taken to imply intentional agency, it in turn provides what we might think of as the *bronze* standard for consciousness:

Bronze Standard: Command following is a marker of consciousness; i.e. evidence for command following provides strong evidence for the presence of consciousness.⁴

Given that the patients at issue in these studies are incapable of verbal report, they do not satisfy the gold standard. Given that they do not exhibit signs of endogenous agency (despite apparent preservation of the neural and motoric machinery required for endogenous agency), they do not satisfy the silver standard. That leaves the bronze standard as the only option for researchers. Reliance on the standard is clearly visible in the work of Owen et al. (2006) and Monti et al. (2010).

In what follows, our aim is to cast doubt on the use of the bronze standard in such studies. While we do not dispute that this standard provides a useful means of studying consciousness in normal subjects, we have serious concerns about its application in severely brain injured patients. Why accept the inference from command following, a capacity to follow task instructions, to the presence of intentional agency, and thence to consciousness?

⁴ Since command following is normally taken to *imply* intentional agency, whereas agency *signifies* consciousness, the link between command following and agency is strictly stronger than the link between agency and the attribution of consciousness.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

Command Following and Akinetic Mutism

Command Following looks rather shaky in the light of the following. As recently discussed by Klein, the responsive PVS patient group bears many similarities to patients suffering from akinetic mutism (AM). AM is a wakeful state in which patients exhibit a severe and persistent decrease in responsiveness; a lack of spontaneous motor or verbal activity; indifference to pain, thirst, or hunger; flattened affect; and apathy without depression. As Klein interprets the data from AM patients, they are not capable of endogenous intention formation, and hence they lack a capacity for endogenous agency.⁵ Despite their lack of spontaneous activity, AM patients are capable of instruction following and answering questions. Given the right prompts, AM patients may even engage in complex activity, such as reading a test and answering questions about it. Following Klein, we classify such responses as *stimulus-evoked cognition*, as contrasted with endogenously initiated intention and action. AM patients are claimed to be acting on commands without the mediation of endogenous intention, and hence without endogenous agency. As such, Command Following simply doesn't stand up in AM.

The PVS patients in the cited studies and AM patients have damage to much the same parts of the brain.⁶ In particular, both patient groups exhibit damage to pre-supplementary motor area (SMA), which is heavily implicated in the literature on AM, and which many studies suggest is associated with voluntary, endogenous, actions. Both patient groups exhibit a complete absence of self-initiated action or

⁵ Klein adopts a scientific notion of 'intention', which 'simply stands for whatever internal motivational state gives rise to a particular action, subject only to the restriction that it is a sufficiently complex state that it cannot occur completely automatically.'

⁶ For more details see Klein §3.3.1.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

response, which is not explained by paralysis. Moreover, the studies of Owen et al. (2006) and Monti et al. (2010) suggest that these PVS patients likewise retain a latent capacity to follow task instructions and even answer questions. Given the structural and behavioural similarities, then, it is plausible to suggest that responsive PVS patients are complying with the experimenters' requests in an analogous (or even identical) manner to the way in which AM patients engage in this behavior.⁷ On this model, responsive PVS patients are exhibiting stimulus-evoked cognition in response to task commands, but not an endogenously generated intention to imagine playing tennis or navigating about one's house. And if this is right, then Command Following doesn't hold up in the study of borderline states of consciousness either.

One immediate response would be that although the command following exhibited by these patients is stimulus evoked and hence not endogenously initiated, it still signals a residual, functioning, capacity for agency in these subjects. Even though these events are exogenously stimulated, the thought goes, they nonetheless establish the presence of an intentional response, and that is all that is required to establish agency, and hence consciousness. While this move is tempting, we think it ought to be resisted. Let's suppose for the sake of argument that the data do establish the presence of an intentional response, by which we mean an event that is initiated and guided by an intention, an internal motivational state that prompts action (including mental action). We think that intentional responses in this sense do not suffice to establish *agency*. Agency is best understood as requiring the ability not simply to act or

⁷ Note that the claim is not that responsive PVS patients *are* AM, just that there is a shared deficit of endogenous intention, and hence of endogenous agency. The explanatory power of such a model should be judged on its merits, then, not on the plausibility of a diagnosis of AM in the PVS patient group.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

respond, but to act in a way that is relatively independent of stimuli, and that requires considerable flexibility of response. What is required is not simply the formation and execution of intentions, but of *endogenously initiated* intentions. The patient groups under discussion do not exhibit agency in this sense. The responsive PVS patients do not (as far as we know) issue responses of this sort without external prompting, and their responses are moreover extremely inflexible and constrained. Even granting the presence of intentions to act or imagine, then, command following does not suffice in this case for the attribution of agency, and hence does not suffice to establish the presence of consciousness.

Others are free to define ‘agency’ as they wish, of course, and one can imagine many less demanding notions than that sketched above. In particular, some may wish to tie agency to mere intentional response, even when that response is stimulus bound and inflexible. However, we think that this attenuated (by our lights) notion of agency provides a weak foundation for the attribution of consciousness to responsive PVS patients. At most, it seems to us to provide grounds for ascribing consciousness to patients *only when they are exhibiting a stimulus evoked intentional response*.

Consciousness here would be ‘fleeting and stimulus-bound’, as Klein puts it.⁸ This is because if agency is the operative marker for consciousness, and if ‘agency’ is exhibited *only* when exogenously evoked, and not at all otherwise, then presumably the evidential situation warrants the attribution of consciousness only when the subject is following the experimenters’ commands. In contrast, agency in our

⁸ Klein considers and rejects this interpretation of ‘fleeting’ consciousness, which has been proposed elsewhere in the literature by Damasio (2000).

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

preferred sense, which requires an endogenously initiated and flexible response, would provide much stronger evidence for a standing or uninterrupted conscious state.

Let's take stock. So far, we've argued that Command Following looks very shaky in the case of AM. Given the neural and behavioural similarities between AM and responsive PVS patients, it is plausible to model responsive PVS patients as sharing a deficit in endogenous intention formation. If this model holds, Command Following also looks shaky in these studies at the borders of consciousness, and this in turn entails that the responses of PVS patients do not provide strong evidence for the presence of. The PVS data can be interpreted as involving nothing more than stimulus-evoked cognition, hence no endogenous agency, and hence no consciousness.

As we'll see in the next section, Klein nonetheless thinks that extreme AM patients and responsive PVS patients are in fact conscious. Given the foregoing arguments, clearly this inference has to go via another route besides Command Following and Agency. Our next task will be to critically evaluate his reasons for attributing consciousness in such cases.

Another Route to Attributing Consciousness?

Despite casting doubt on Command Following, Klein nevertheless suggests that AM patients are conscious in some sense. It follows that he also sees no reason to deny consciousness to responsive PVS patients. In what sense are these patients deemed to

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

be conscious? And on what grounds, if not via the Bronze Standard, are we supposed to attribute consciousness in such cases?

Regarding the first question as to what type of consciousness is at issue, Klein suggests that AM patients possess *peripheral* consciousness without *focal* consciousness, a distinction derived from Kriegel (2004). As we understand it from Kriegel, peripheral consciousness is the sort of awareness that one has, for example, of objects in peripheral vision, the sound of passing cars as one listens to a concerto, or a background sense of cheerfulness resulting from my good mood. Focal consciousness in contrast is the sort of awareness one has of a foveated section of text held in front of one's face, or the melody line played by the soloist of a concerto. Klein seems to use 'peripheral consciousness' in a somewhat looser sense: he views focal consciousness as involving the allocation of *attention*, whereas peripheral consciousness involves some awareness of *unattended* stimuli.⁹ We think that nothing crucial hangs on this issue in what follows.

As for the second question, regarding the alternative route to attributing consciousness, Klein draws on two sources of evidence for the claim that AM patients are conscious (in the peripheral sense). The first source is self-reports by symptomatic AM patients. Patients with less severe forms of AM report a curiously 'empty' mental state; these introspective reports are taken by Klein to be extremely good evidence that they are conscious. The second source concerns retrospective reports by patients who have recovered from AM. In some (but, importantly, not all) cases, they report

⁹ Klein has endorsed our interpretation in correspondence.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

having had some consciousness of their surroundings and of events when they were symptomatic.

We are skeptical that either source of evidence provides strong grounds for attributing consciousness to AM patients. Let's first consider self-reports by symptomatic AMs. Here we should distinguish between two possible cases. The first is self-reports in patients with mild to moderate AM, who we assume are capable of *some* degree of spontaneous, self-initiated, activity. The second is self-reports in patients with extreme cases of AM, who let's assume are completely incapable of any degree of spontaneous, self-initiated, activity.

There are two worries concerning the first type of case. First, it is unclear why self-reports in such moderate cases of AM should be considered relevant to the argument concerning consciousness in PVS patients. PVS patients, we assume, are most similar to extreme AM patients, if indeed they are similar to AM patients at all. They are most akin to extreme AM patients because (like them) PVS patients exhibit no spontaneous, self-initiated, activity whatsoever. As such, it is of questionable relevance that some moderate AM patients, who are capable of some self-initiated activity, self-report states of consciousness. A second, related, worry is that if moderate AM patients are capable of some spontaneous activity, it is reasonable to infer that they retain some capacity for endogenous intention and agency. But if it is antecedently known that moderate AM patients have some capacity for endogenous agency, then self-reports of consciousness are not even essential: we can simply appeal to Agency, the claim that intentional agency is a strong marker for the

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

presence of consciousness. (Recall that our target above was Command Following, not Agency: for all we've said, Agency is just fine.) For moderate AM patients, then, we have an alternative, independent, basis for attributing consciousness – a basis that is unavailable in PVS patients, for whom spontaneous activity, hence clear signs of endogenous agency, are crucially lacking. This disparity makes it wholly illegitimate to appeal to putative examples of consciousness in moderate AM patients, in order to support claims of consciousness in responsive PVS patients.

We can also raise concerns regarding the second type of case, that is, self-reports in extreme AMs.¹⁰ As we've characterized them, extreme AM patients are completely incapable of any degree of spontaneous, self-initiated, activity. As such, we can assume that any self-report issued by an extreme AM would have to be prompted or otherwise elicited. This creates a highly unusual situation regarding the evidential status of such reports. In order to see this point, let's briefly rehearse the dialectic to this point. Our overarching question is whether extreme AM patients, likewise responsive PVS patients, are conscious. Both we and Klein agree that mere responsiveness to commands, or stimulus-evoked cognition, is insufficient to establish the presence of consciousness. Klein nonetheless believes that AM patients *are* conscious, and is therefore required to provide some independent grounds for attributing consciousness in such cases. And now here's the puzzle: what are we to say when the tendered grounds for attributing consciousness are stimulus-evoked self-reports regarding the patient's putative state of consciousness? Our inclination is to

¹⁰ We're not in fact sure whether any such reports exist, but this doesn't matter. The argument goes through all the same.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

say that if stimulus-evoked cognition *in general* provides insufficient grounds to attribute consciousness, then we should also conclude that stimulus-evoked self-reports of consciousness provide insufficient grounds to attribute consciousness. We admit that there might be some superficial weirdness in this response, but it strikes us as the logical conclusion given the foregoing discussion.

We've been discussing Klein's first source of evidence for consciousness in AM self-report in symptomatic patients. Klein's second source of evidence is retrospective reports by patients who recover from AM. Some of the problems raised above for self-report also apply to retrospective reports. In particular, if the retrospective reports are from someone who has recovered from mild or moderate AM, then they are of dubious relevance to the issue at hand. In order to focus on the most challenging case, however, let's consider a retrospective report from someone seemingly recovered from extreme AM.¹¹ Klein discusses Laplane's (1984) patient, who describes his past state as like having 'a blank in my mind'. When quizzed on private thoughts, Bogousslavsky's et al.'s (1991) patient said 'I think of nothing', 'I don't want anything'. Another said they 'did not have any projects for the future and did not have any personal thoughts'.

While extremely interesting, it is unclear that statements of this sort provide any real evidence regarding the putatively conscious states of the individual while they were symptomatic. First, there are general worries about the quality and accuracy of

¹¹ Again we are not sure whether such retrospective reports from genuine *extreme* AMs exist. This does not matter for our argument.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

testimony in such cases. Given the strikingly bizarre, uncommon, and elusive nature of extreme AM, perhaps it is simply not reasonable to expect recovered subjects to have well formed, reliable, thoughts about their symptomatic state. Second, all of the reports quoted above concern absences or failures, such as the absence of thought or desire. It is extremely unclear how to interpret such reported absences. One option would be to infer that the patient was, while symptomatic, in some sense conscious of having a ‘blank in their mind’, and that this memory is reported in their retrospective appraisal. Another option, which seems at least equally plausible, however, is that the patient is just trying to find a way to report an absence of consciousness. These retrospective attempts might be aided or enhanced by the fact that the subject, while symptomatic, *ex hypothesi* would have been capable of certain types of stimulus-evoked cognition, and hence would have retained some degree of informational sensitivity to sensory input. Recovered patients might even have access to some of the information that was taken up by the cognitive system during this period. But recalling informational content that might have been conscious in a normal subject is not necessarily recalling consciousness. Recovered patients might also acquire some sort of awareness that their explicit beliefs, desires, and intentions did not change during the symptomatic period, and might report this by saying that they ‘wanted nothing’ or did not ‘have any projects’ during this period. Such retrospective appraisals, however, clearly do not force the attribution of conscious awareness of such absences while they were symptomatic.

Further analysis and careful reflection on these puzzling cases is certainly required.

As things stand, however, we find no compelling reason to think that extreme AM

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

patients – the most relevant group, as far as the comparison with PVS patients is concerned – are conscious. At the very least, if there *is* any consciousness in such patients, these states most likely would be fleeting and stimulus-bound.

The Moral Status Question

In this final section, we consider the moral significance of purported attributions of consciousness in PVS patients. Notwithstanding the skeptical conclusions of the previous two sections, let's suppose for the sake of argument that Klein is correct and that extreme AM patients are conscious, and that responsive PVS patients are similar enough in terms of their deficits to warrant the attribution of consciousness to them too. Let's even suppose that both sets of patients are conscious all (or much) of the time, and not just when they are prompted. Would they then enjoy the moral status that is rightly attributed to normal subjects in virtue of the fact that they are conscious? We suggest that the answer is no.

One of us has previously argued that the bulk of the work in underwriting moral status is done not by the capacity for experience per se, but by the capacity to have series of appropriately linked mental states (Levy 2009; Levy & Savulescu 2009). A being with an interest in having a life must be able to care how its life goes, and this requires very sophisticated cognitive abilities, such as an ability to conceive of oneself as a being persisting through time, to recall one's past, to plan and to have preferences for how one's life goes (Singer 1993; McMahan 2002).

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

It seems clear that phenomenal consciousness is neither (metaphysically) necessary nor sufficient to sustain this kind of interest in having a life. Rather, it seems that what is required is for the being to self-attribute informational states of the right kind. This in turn requires the simultaneous availability of these informational states to (some of) the consuming systems constitutive of its mind. It follows that a necessary condition of possessing a serious interest in life is something akin to Block's (1995) notion of access consciousness, though – because we wish to avoid building in some of the commitments of Block's notion, such as the stipulation that information of which a being is access conscious is poised for the rational guidance of its behavior – we shall call this kind of consciousness *informational consciousness*.

To have a serious interest in life, a being must self-attribute an appropriate set of its mental representations, where the content of these representations concern its own plans and projects, and the ground of these plans and projects: its own existence. A serious interest in life therefore requires self-consciousness, of this self-attributing kind, which in turn depends on informational consciousness in the sense outlined above. Only a being that is self-conscious of an appropriate set of representations can hold itself as having plans and projects and therefore possess future-oriented desires with regard to these plans and projects. Only a being like this can suffer the distinctive harm associated with death: the permanent thwarting of these plans and projects. Only a being who can suffer such a harm has a reason to fear death (rather than dying).

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

Interpreting the moral significance of the foregoing experimental results therefore requires answering several questions, empirical and conceptual, in order to establish whether patients possess the capacity for informational consciousness (with the right contents) and for self-consciousness (again, with the right contents). As we saw, Klein attributes *peripheral consciousness* to the PVS patient group. We have expressed doubts about this attribution, but for the purposes of discussion we accept it. In the previous section, peripheral consciousness was characterized as the sort of awareness one has of items in the periphery of one's visual field, for example, or of a background mood of happiness. Perhaps this type of consciousness accompanies objects or states that are *unattended* by the subject. Given this understanding, one way to interpret states of peripheral consciousness is as a subset of states of phenomenal consciousness. The thought here is that if one's entire visual field constitutes one's total state of visual phenomenal consciousness, then presumably points in the periphery of this field would constitute a subset of states of visual phenomenal consciousness. As pointed out above, however, an exclusively phenomenal notion of consciousness such as this cannot do the work of underwriting a serious interest in life. On this reading, then, the attribution of peripheral consciousness gets us no closer to moral significance than the attribution of purely phenomenal consciousness.

On another interpretation of peripheral consciousness, this notion carries with it implications not, or not just, of phenomenal consciousness, but of informational consciousness in the sense outlined above. On this understanding, one's peripheral consciousness of a background mood of happiness, say, would have informational content that is available to cognitive consuming systems. Even in this case, however,

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

the prospects of underwriting a serious interest in life are not much improved. We assume that states of peripheral consciousness would have relatively thin, impoverished, indeterminate, and perhaps highly determinable informational contents. Peripheral consciousness therefore would offer very little to one's consuming systems in the way of specific, categorized, or recognitional information about one's environment, one's moods and desires, and so on. And this seems to fall well short of the sorts of content required for the formation, monitoring, and self-attribution of plans and projects by the subject. How could I so much as form, let alone self-attribute, a plan to reach for the glass by my bedside, if my awareness of the glass is limited to the sort of vague, indeterminate, uncategorized awareness that one has of objects in the periphery of vision?

There are many complications here that warrant more extended discussion. One such complication is that insofar as responsive PVS patients can be attributed *some* kind of consciousness, they would admittedly have a degree of moral status that completely nonconscious beings lack. More specifically, if these patients possess peripheral consciousness, we assume that they could have some peripheral awareness of pain and pleasure. We are somewhat unclear on what the nature of such peripheral pains and pleasures would be like: for example, would the attention grabbing nature of pain put pressure on the idea that one could be *merely* peripherally conscious of pain? Whatever the truth of the matter, the mere peripheral awareness of pain and pleasure might create obligations to take these states into account, for example, in issuing analgesics or creating environments in which sensory pleasures can be increased.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

A second complication concerns the suggestion (made by Klein) that subjects who possesses a *standing* state of peripheral consciousness might come to have states of focal consciousness when prompted. That is to say, perhaps issuing a command to an unresponsive patient could shift them from an indeterminate and inattentive state of awareness to a state with richer content through greater allocation of attention. If this were the case, it would be reasonable to think that they could come to have focal consciousness given other kinds of stimuli. In particular, they might have full-blown focal awareness of pain and pleasure, states that are plausibly of direct moral significance. Here the moral obligations to take these states into account in making treatment decisions, for example, would be much clearer-cut than in the case of merely peripheral pain and pleasure. Even granting these points, however, one of us thinks that these experiences in normal subjects get a great deal of their moral significance from their wider role in agents' mental lives, not from their raw phenomenal feel (Levy 2009). Unless it could also be established that these patients are also deploying the resources of informational and self-consciousness in registering and self-ascribing these states, then, it follows that pain is not as bad and pleasure not as good for these patients.

Conclusions

The mental lives of patients in extreme and prolonged states of unresponsiveness remains one of the most puzzling and ethically significant issues in contemporary brain science and medicine. We believe that the waves of excitement surrounding putative findings of consciousness in PVS patients have been premature. The methodology encapsulated in the bronze standard, which takes the ability to follow

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

commands as a proxy for consciousness, is of dubious standing. Following a plausible model of the impairments characterizing AM, command following cannot be adopted as a straightforward marker of endogenous intentional agency in such patients. The mere capacity for stimulus-evoked cognition and response in our view does not suffice for agency, and hence cannot provide the desired link to consciousness. Much further work is clearly needed to understand the range and limits of such exogenously initiated intention and action. Such future work will shed further light on the mental states of patients suffering these profound impairments in the capacity for endogenous action. Perhaps at some later stage another route to the attribution of consciousness will be found in such patients. We would welcome such developments. As we have argued, however, even granting such claims of consciousness, their ethical significance should not be immediately inflated. One plausible candidate for the putative consciousness exhibited in such cases, peripheral consciousness, lacks the properties required to ground serious self-interest and hence morally significant agency.

References

Bayne, T. 2013. Agency as a marker of consciousness. In T. Vierkant, J. Kiverstein and A. Clark (Eds.) *Decomposing the Will* (pp.160-80). Oxford: Oxford University Press.

Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18: 227-287.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

Bogousslavsky J, Regli F, Delaloye B, Delaloye-Bischof A, Assal G, Uske A. 1991. Loss of Psychic Self-activation with Bithalamic Infarction. *Acta Neurologica Scandinavica* 83: 309-16.

Giacino JT, Ashwal S, Childs N, Cranford R, Jennett B, Katz DI et al. 2002. The minimally conscious state: definition and diagnostic criteria. *Neurology* 58: 349–53.

Damasio A. 2000. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. New York: Harcourt, Brace, and World, Inc.

Klein, C. forthcoming. Consciousness, Intention, and Command-Following in the Vegetative State. *British Journal for the Philosophy of Science*.
doi: 10.1093/bjps/axv012

Kriegel U. 2004. The Functional Role of Consciousness: A Phenomenological Approach. *Phenomenology and the Cognitive Sciences* 3: 171-93.

Laplane D, Baulac M, Widlöcher D, Dubois B. 1984. Pure Psychic Akinesia with Bilateral Lesions of Basal Ganglia. *Journal of Neurology, Neurosurgery, and Psychiatry* 47: 377-85.

Levy, N. 2009. What Difference does Consciousness Make? *Monash Bioethics Review* 28: 1-13.

Forthcoming in *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*, Walter Sinnott-Armstrong (ed.), New York: Oxford University Press. Please cite published version.

Levy, N. & Savulescu, J. 2009. The Moral Significance of Phenomenal Consciousness. *Progress in Brain Research* 117: 361-370.

McMahan, J. 2002. *The Ethics of Killing: Problems at the Margins of Life*. Oxford: Oxford University Press.

Monti, M. M., Vanhaudenhuyse, A., Coleman, M.R., Boly, M., Pickard, J.D., Tshibanda, L., Owen, A.M. & Laureys, S. 2010. Willful Modulation of Brain Activity in Disorders of Consciousness, *New England Journal of Medicine*, 362: 579–89.

Owen, A.M., Coleman, M.R., Boly, M, Davis, M.H., Laureys, S. & Pickard, J.D. 2006. Detecting Awareness in the Vegetative State. *Science* 5792: 1402.

Singer, P. 1993. *Practical Ethics*. Cambridge: Cambridge University Press.