



# A Critical Notice on the Moral Grounding Question in David Chalmers' *Reality+*

Anand Jayprakash Vaidya<sup>1</sup>

Accepted: 21 December 2022

© The Author(s), under exclusive licence to Springer Nature B.V. 2023

## Abstract

In this critical discussion, I evaluate David Chalmers' position on the moral grounding question from his (2022) *Reality+*. The moral grounding question asks: in virtue of what does an entity  $x$  have moral standing? Chalmers argues for the claim that phenomenal consciousness is a necessary condition for moral standing. After a brief introduction to his book, I evaluate his position on the moral grounding question from the perspective of access consciousness as opposed to phenomenal consciousness, as well as the Jain doctrine of non-violence, and the differentiation of creatures in terms of their sense capacities.

## Introduction: the Book

In *Reality+*, David Chalmers, argues for the thesis that *virtual reality is genuine reality*. His book engages three big questions: the knowledge question, the reality question, and the value question.

*The knowledge question*: can we know whether or not we are living in a virtual world? In part 2, Chalmers argues that we cannot know that *we are not* living in a virtual world. He does so by defending his own version of the simulation argument, which *very roughly* goes as follows: (1) If there is nothing (technological or otherwise) that would serve to block the creation of a perfect simulation of a humanlike being into existence (what he calls a sim blocker), then most humanlike beings are simulations. (2) If most humanlike beings are simulations, we are probably simulations. So, (3) if there are no sim blockers, we are probably sims.

*The reality question*: are virtual worlds real? Chalmers argues that it isn't improbable that we are living in a simulation. However, it doesn't follow that the objects that inhabit our simulation are real. Thus, in chapter 6 Chalmers offers five

---

✉ Anand Jayprakash Vaidya  
Anand.Vaidya@sjsu.edu

<sup>1</sup> Department of Philosophy, San Jose State University, 1 Washington Square, San Jose, CA 95192, USA

definitions of “reality” and shows that virtual worlds satisfy these definitions in the same sense that non-virtual entities do.

Let C be any ordinary object you take to be real, such as a cup. Let V be any virtual object you take to be non-real, such as a character in a video game. Chalmers considers multiple definitions of reality and shows that, just as C exists, has causal powers, is mind-independent, non-illusory, and genuine, so is V. V exists in the video game, it has causal powers that players of the game interact with, it doesn't depend on the mind of the player but on the video game, it isn't an illusion, and it is genuine (not a counterfeit). So, if C is real for those reasons, so is V.

The *value question*: can you lead a good life in a virtual world? In parts 5 and 6, Chalmers argues that one can. He holds that in principle, life in virtual reality can have the same sort of value as life in non-virtual reality. That is, life in virtual reality can be good or bad. In answering the value question, Chalmers also takes on the moral grounding question: in virtue of what does something have moral status? It is this question and his answer to it that I found most interesting. I now wish to critically examine the answer.

## The Moral Grounding Question

Some 40 years ago, Ned Block carefully distinguished between two types of consciousness. P-consciousness, phenomenal consciousness, captures what Thomas Nagel called the ‘what it is like’ aspect of subjective experience. There is something it is like to see blue vs. red, smell curry vs. marinara sauce, hear C minor vs. F minor, touch sand vs. water. A-consciousness, access consciousness, by contrast, is about capacities, such as reasoning, using language, and engaging in rationally guided action. These two definitions of consciousness can come apart. For example, Block discusses blindsight patients who report that they cannot see anything in their blind field, but can do lots of things, such as grasp objects and navigate, through the blind field. Access consciousness is much more closely tied to the kind of intelligence that Alan Turing was investigating in his classic paper on computing machinery and intelligence, phenomenal consciousness is not. And although Block did not distinguish between types of phenomenally conscious states, one can distinguish, as Chalmers does, between various types of phenomenal consciousness. One subclass of phenomenal consciousness is affective consciousness. Good examples of affectively conscious states are love, anger, joy, and sadness. On a non-cognitivist account of what an emotion is, emotions are states of affective consciousness.

In response to the moral grounding question, Chalmers argues that phenomenal consciousness is necessary for moral status. He holds that affective consciousness is not and does not take a stand on access consciousness. To argue that phenomenal consciousness is necessary for moral status, Chalmers makes use of thought experiments that build off of Phillipa Foot's famous runaway trolley problem. Chalmers does this by considering non-human creatures called ‘zombies’ and ‘vulcans’ within the framework of the trolley problem. He then argues that the capacity for phenomenal consciousness as opposed to the capacity for affective consciousness is necessary for moral status.

In the original runaway trolley problem, one is asked whether one should change the direction of a runaway trolley by moving a switch so as to direct the trolley toward killing one person and away from running over five. If you switch tracks, you kill one person, but if you don't switch tracks, you let five people die. Many have the reaction that one should kill one person in order to save five. To be consistent with the original thought experiment, but get at the spirit of Chalmers' argument, I have reconstructed his alteration of Foot as follows.

*Zombie trolley problem:* Zombies are physical duplicates of humans, but lack phenomenal consciousness and by consequence lack affective consciousness. There is nothing it is like to be a zombie. You can change the path of a runaway trolley by moving a switch. If you do nothing, the runaway trolley will kill five zombies who are on the tracks in front of it. If you switch the tracks, you will kill one human, but save five zombies. What should you do?

*Vulcan trolley problem:* Vulcans are physical and phenomenally conscious duplicates of humans, but they lack states of affective consciousness, such as joy, happiness, sadness, and anger. There is something it is like to be a vulcan, because vulcans are phenomenally conscious. They only lack affective consciousness. You can change the path of a runaway trolley by moving a switch. If you do nothing, the runaway trolley will kill five vulcans who are on the tracks in front of it. If you switch the tracks, you will kill one human, but save five vulcans. What should you do?

Chalmers' view is that the capacity for phenomenal consciousness, and not the capacity for affective consciousness, is necessary for moral status. It emerges, at least, in his remarks about what holds in the zombie vs. the vulcan trolley problem:

If you think a single conscious creature should be saved at the cost of killing five non-conscious, this suggests that consciousness is *relevant* for moral status. Conscious creatures matter more than non-conscious creatures. If you hold the stronger view—that there's never a moral reason to spare non-conscious creatures—this suggests that consciousness is *necessary* for moral status. Non-conscious creatures don't matter at all, morally speaking. The stronger conclusion dovetails with the view I advocated at the end of [Chapter 17], that consciousness is the ground of all value. (Chalmers, 2022: 341).

And:

My own sense is that a Vulcan matters about as much as an ordinary human. Of course I am glad that I am a human and not a Vulcan, since affect makes my life better. Suffering and happiness make a big difference to how good or bad a conscious creature's life is. *But they're not what gives a creature moral status.* (Chalmers, 2022: 344, *emphasis added*).

Chalmers is arguing that while it is permissible to let five zombies die, rather than kill one human, it is impermissible to let five vulcans die so as to avoid killing one human. The capacity for affective consciousness is not a necessary condition on

moral status. However, the capacity for phenomenal consciousness is necessary. So, just as you should kill one human to save five humans, you should kill one human to save five vulcans.

Often philosophers, such as Peter Singer, are motivated to hold that a creature has moral status because they have the capacity to physically suffer. For example, not only can humans physically suffer, but so can apes, octopi, dolphins, and likely even lobsters. However, in pain asymbolia, a human fails to be able to feel pain even though they can have damage to their body. In such a case, a person fails to have the key component that is the ground of moral standing: the capacity to suffer physically. Thus, humans with pain asymbolia are a challenge to the capacity to suffer account of the ground of moral status.

Chalmers can avoid this critique because he thinks that phenomenal consciousness and not the capacity to suffer is the necessary condition for moral standing. A person with pain asymbolia is phenomenally conscious; they lack the capacity for phenomenal suffering. As a consequence, Chalmers faces a distinct objection. One that arises from cases where humans are in an irreversible coma. Patients in an irreversible coma don't have phenomenal consciousness; as a consequence, they don't have moral standing. So, if we want to extend moral standing to them, we need to work around phenomenal consciousness.

Both Singer and Chalmers could respond to the objections from a person with pain asymbolia and a person in an irreversible coma by giving the *species membership* response. Humans who either have pain asymbolia or are in an irreversible coma are still members of a species that generally has the capacity for pain and the capacity for phenomenal consciousness. For some this would not be satisfying because the attempt to identify the capacity for pain and phenomenal consciousness as the ground of moral status aims to move around referencing species membership. Peter Singer, for example, made famous the expanding circle argument for moral standing by arguing that non-human animals can feel pain just like humans and thus it is inconsistent to grant moral standing only to humans and not to non-human animals. The aim is to find a multiply realizable property across various kinds of entities, and not a species specific property.

The challenges from pain asymbolia and coma patients leads me to the conclusion that neither the capacity to feel pain nor the capacity for phenomenal consciousness is necessary for moral standing. Rather, they are only sufficient for moral standing. A creature can have moral standing in virtue of either having the capacity for phenomenal consciousness or affective consciousness (especially, the capacity for phenomenal suffering). But these are not necessary. In my piece, 'If a robot has consciousness, is it okay to turn it off?' I consider a future where robots exhibit the capacity for access consciousness, but not phenomenal consciousness; I argue that access consciousness is also sufficient for moral standing. Currently, there are no machines that have access consciousness. However, in the development of AI, we should ask: what does it take for a machine to have access consciousness?

According to Montemayor and Mindt (2020), artificial general intelligence (AGI) is necessary for access consciousness because of the way in which AGI yields a

form of attention that makes information poised for use. By AGI, I mean the capacity to exhibit rationally guided action across multiple domains. For example, current AIs are remarkable in that they can routinely beat the best humans at Chess or Go, or solve a protein folding problem. But, to date, machines cannot do multiple tasks intelligently. No single machine can produce music, bake a cake, and fold clothes into a suitcase. However, on my view there are no in-principle barriers to the development of AGI in machines. It is not impossible for AIs to have AGI.

Thus, in opposition to Chalmers and Singer, my view is nicely summarized by the idea that clusters of capacities are sufficient for moral standing. Sure, one can have moral standing because one can feel pain, and one can have moral standing because one has the capacity for phenomenal consciousness. But one can also have moral standing because one has certain general capacities and abilities, which are distinct and independent from the capacity to feel pain and be phenomenally conscious. Machines, for example, could have moral standing because they can have access consciousness but not phenomenal consciousness. In addition, on the view I offer, it is possible for a person who has pain asymbolia or a person that is in an irreversible coma to have moral standing not because they belong to a specific species but because they have some other capacities that matter. Granted Chalmers is correct that more work needs to be done for the kind of view I advocate to be taken seriously because unlike phenomenal consciousness, which has pain – which we all seem to care about avoiding, it is unclear why access consciousness or any other capacity that doesn't depend on phenomenal consciousness would matter.<sup>1</sup>

For a comprehensive evaluation of Chalmers work it is important to draw it into conversation with non-western views of moral standing. One tradition to draw it into conversation with is Jainism. Jains offer an expansive account of moral status. Jains are the original proponents of *ahimsā*, which holds that violence is to be minimized as much as possible. From the perspective of Jainism, everything from microbes to humans have a soul (*jīva*). Having a soul means you morally matter. Harm to souls is to be avoided and minimized by humans. Two key ideas are (i) the distinction between the number of sense organs and sense capacities that various creatures have, and (ii) the overall view that harm should be minimized. (i) is central to practical debates about moral standing because it is almost impossible to live without harming, and thus we need a graded view of moral standing, one that offers a reason as to why certain entities have more or less moral standing than others.

Jains don't distinguish between phenomenal consciousness and affective consciousness. Rather, they distinguish between various kinds of sense capacities that an ensouled creature has, given its embodiment. For example, one-sense

---

<sup>1</sup> See Bradford (2022) for a great discussion of whether or not phenomenal consciousness is necessary for moral standing and for being a welfare subject. Bradford's discussion challenges the idea that appealing to phenomenal consciousness explains why something is a welfare subject. Here I have assumed that those that appeal to phenomenal consciousness as the ground for moral standing have the advantage, however, there are those, like myself, who question that position.

creatures have a soul but can only touch, two-sense creatures have a soul but can touch and taste, while five-sense creatures have a soul that can touch, taste, see, hear, and smell. While the taxonomy might no longer be scientifically defensible, it is the functional role of the taxonomy that is interesting and survives as a theory worth engaging with. The number of senses plays a role in determining how much suffering the creature can undergo given its embodiment. The capacity to sense is what leads to a gradient view of moral status. For example, two-sense creatures might change direction of movement in reaction to the amount of oxygen or salt in their environment, but they do not have the capacity for suffering sadness. Five-sense creatures, such as mammals, have a capacity for greater suffering, as evidenced by the anxiety shown at the loss of offspring both in humans and in bovine. Jains hold that given the embodiment of a soul in a body plan with a certain set of sensory capacities, a gradient view of moral standing is generated. They endorse a version of the thesis that *when x has more F than y, x has more moral status than y*. *F* on their view is sense capacities. On another theory *F* might be phenomenal consciousness, where the view would be that *when x has more consciousness than y, x has more moral status than y*.

While the Jain account of the gradient of pain and suffering can be scientifically challenged, it is nevertheless relevant because it allows us to answer applied questions in research ethics on the use of human and non-human animals in the production of knowledge. There are three Rs in research ethics: replacement, reduction, and refinement. Replacement says we should try to replace research on humans with research on non-human animals, when doing so would not affect the quality of the knowledge gained. Reduction says we should try to reduce the number of animals we test on when doing so would not affect the quality of the knowledge gained. Refinement says that we should try to do research on animals that feel less pain, or pain less intensely, when doing so would not affect the quality of the knowledge gained. The Jain theory offers an answer to which animals feel more pain and why. This gradient view of pain across species feeds nicely into the actual rules for research on human and non-human animals.

## References

- Bradford, G. (2022). Consciousness and welfare subjectivity. *Nous*, 1–17. <https://doi.org/10.1111/nous.12434>
- Chalmers, D. (2022). *Reality+: Virtual worlds and the problems of philosophy*. Norton & Company Inc.
- Montemayor, C., & Mindt, G. (2020). A roadmap of artificial general intelligence. *Mind & Matter*, 18(1), 9–37.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.