



# Borderline consciousness, when it's neither determinately true nor determinately false that experience is present

Eric Schwitzgebel<sup>1</sup>

Accepted: 29 August 2023 / Published online: 4 October 2023  
© The Author(s) 2023

## Abstract

This article defends the existence of *borderline consciousness*. In borderline consciousness, conscious experience is neither determinately present nor determinately absent, but rather somewhere between. The argument in brief is this. In considering what types of systems are conscious, we face a quadrilemma. Either nothing is conscious, or everything is conscious, or there's a sharp boundary across the apparent continuum between conscious systems and nonconscious ones, or consciousness is a vague property admitting indeterminate cases. Assuming mainstream naturalism about consciousness, we ought to reject the first three options, which forces us to the fourth, indeterminacy. Standard objections to the existence of borderline consciousness turn on the inconceivability of borderline cases. However, borderline cases are only inconceivable by an inappropriately demanding standard of conceivability. I conclude with some plausible cases and applications.

**Keywords** Consciousness · Phenomenology · Phenomenal consciousness · Vagueness · Indeterminacy · Animal consciousness · Semi-conscious states · Minimal consciousness

## 1 Introduction: the central claim, some clarifications, and the argument in brief

I will defend the existence of *borderline consciousness*. In borderline consciousness, conscious experience is neither determinately present nor determinately absent, but rather somewhere between. It is neither determinately true nor determinately false that there's something it's like to be the entity in question or to have the state in ques-

---

✉ Eric Schwitzgebel  
eschwitz@ucr.edu

<sup>1</sup> Department of Philosophy, University of California, Riverside, Riverside, CA 92521-0201, USA

tion. The existence of such indeterminate gray zones of consciousness has sometimes been defended on general evolutionary or metaphysical grounds or in view of seemingly “twilight” or transitional states of consciousness, coming in or out of sleep, anesthesia, or brain injury.<sup>1</sup> Others have rejected the possibility of borderline cases, often on the grounds that they are inconceivable.<sup>2</sup> Rarely, however, have arguments in favor been systematically presented or arguments against systematically rebutted.<sup>3</sup>

By “consciousness” I mean *phenomenal consciousness* as the term is standardly used in recent Anglophone philosophy. The term is best defined by example. Look around a bit, considering your visual sensations as you do so. Pinch the skin on the back of one hand, observing the mild pain. Contemplate who you’d revive from the dead if you could, and notice the thoughts or images at the forefront of your mind. The visual sensations, the felt pain, and the thoughts or images all share an obvious property in common. They are all *conscious experiences*. There’s “something it’s like” to have them (in Thomas Nagel’s (1974) memorable phrasing). They have qualitative character. This obvious property is, on standard views, absent from some other mental states or processes, such as your knowledge (not actively recalled until just now) that Obama was U.S. President in 2010, or the processes by which your visual system converts retinal input into experienced shapes, or the fine-grained procedures keeping you in balance while walking. A conscious state or process is a state or process with the property that is most obviously present in the first set of cases and absent (or assumed to be absent) in the second set.<sup>4</sup> This article defends the position that there are borderline cases of that property.

Phenomenal consciousness is thus a *vague* property. It’s like the property of being bald, or green, or extraverted. It’s unlike the property of being exactly equal to 4 or the property of a hydrogen electron’s being in the lowest “ground state” orbital. Baldness, greenness, and extraversion admit of indeterminate, borderline, in-betweenish cases: a man in the gray zone just shy of outright baldness, a turquoise that’s kind of blue and kind of green but not straightforwardly one or the other, a semi-extravert. We can imagine spectra of cases from not-bald to bald, from not-green to green, from not-extravert to extravert. For such spectra, it seems wrong to point to the loss of a single hair, or a minuscule just-noticeable-difference between this shade and that, or the tiniest shred more readiness to say “yes” to party invitations and say *there! That’s exactly the place!* where non-baldness, non-greenness, or non-extraversion become baldness, greenness, or extraversion. In contrast, a number either is exactly equal to 4 or it’s not, and (disregarding superposition) a hydrogen electron is either determinately in its ground state orbital or determinately not, with no borderline cases. Vague

<sup>1</sup> For example, Unger, 1988; Tye, 1996; Dennett, 1998; Overgaard & Overgaard, 2010; Bruno et al., 2011; Casali et al., 2013; Godfrey-Smith, 2020; Kammerer, 2022.

<sup>2</sup> For example, McGinn, 1982/1996; Campbell, 1984; Searle, 1992; Strawson, 1994; Chalmers, 1996; Antony, 2008; Goff, 2013; Bayne et al., 2016; Simon, 2017; Carruthers, 2019; Roelofs, 2019.

<sup>3</sup> Exceptions include Papineau, 1993, 2002, 2003, and Tye, 2021, whose views will be contrasted with my own, in Sects. 8 and 10; Chin, 2015, who defends a view similar to Papineau’s in a PhD dissertation; and Brogaard, in an unpublished manuscript dated 2010, whose argument is primarily grounded in facts about ordinary language use and thus different from the arguments here.

<sup>4</sup> For a fuller discussion of the virtues of defining consciousness by example and an assessment of the risks of doing so, see Schwitzgebel, 2016.

properties – including on the view I will defend the property of being phenomenally conscious – have “gray zones” inhabited by “in-between” or “borderline” cases in which the property is neither determinately present nor determinately absent. The logic of vagueness is contentious, however. If you prefer a different approach, I hope you can suitably translate my claims.<sup>5</sup>

My claim is not merely that conscious experiences can have vague *contents*. You might be able to visually imagine a speckled hen without visually imagining how many speckles it has. You might see a string of text in peripheral vision without seeing exactly what words compose the text. If someone’s name is on the “tip of your tongue” you might, without recalling the exact name, have a rough sense of the name’s approximate length, its first letter, and whether it’s common or rare. You might determinately have such experiences, despite some indeterminacy in their contents. It determinately feels like something to imagine that hen. That’s not borderline consciousness. Borderline consciousness is more difficult to imagine or remember. In fact, as I’ll argue in Sect. 7, there’s a sense in which borderline consciousness is impossible to imagine or remember.

My argument in brief is this. In considering what types of systems are conscious, we face a quadrilemma. Either nothing is conscious, or everything is conscious, or there’s a sharp boundary across the apparent continuum between conscious systems and nonconscious ones, or consciousness is a vague property.<sup>6</sup> Assuming that we accept some form of mainstream naturalism about consciousness, we ought to reject the first three options, which forces us to the fourth, vagueness. Close examination reveals that we ought to be unmoved by standard objections to treating consciousness as a vague property with borderline cases. I’ll conclude with some plausible cases and applications.

## 2 What systems are conscious? A quadrilemma

Are frogs conscious? Garden snails? Lizards, earthworms, honeybees, jellyfish, sea sponges, housecats, black widow spiders, trout, coral? Might some plants be conscious? Under what conditions, if any, might an artificial system such as a computer

<sup>5</sup> Admittedly, some of my arguments below might not translate easily for epistemicists about vagueness such as Williamson, 1994. To those who deny indeterminacy I offer this start of a translation: Reconstruct the fundamental question as the question of whether “conscious” is more like “green” or more like “being in the ground state orbital”. If the arguments below are correct, it’s more like “green”. Then apply your favorite treatment about why it’s hard to find, and why we don’t usually care much about, the exact edges of “green”. Hopefully, the arguments of this article work without such need of translation for most other accounts of vagueness, such as supervaluationism, subvaluationism, contextualism, and theories with more than two truth values, as long as they treat borderline cases as indeterminate or indefinite (for reviews, see Sorensen, 1997/2022; Salles, 2021). I don’t intend to exclude other sources of indeterminacy in statements about phenomenal consciousness, including possibly indeterminacy due to ambiguity (see Sect. 8), incompleteness of a fiction or model, an open future, problematic presuppositions, or under-specified conditionals.

<sup>6</sup> Compare Goff, 2013. Goff finds vagueness unacceptable on broadly the types of inconceivability grounds discussed in Sects. 7 and 9 and thus embraces panpsychism instead.

be conscious? Even without attempting to settle on a specific set of answers, we can see four general shapes that a set of answers might take.

1. *Saltation*. There's a line in the sand. Some systems are determinately conscious and others are determinately nonconscious, with no gray area. Among animals, for example, there might be one or more minimally conscious species that just barely (in their better moments?) possess consciousness, while other species, virtually identical but ever so slightly less sophisticated, lack consciousness. Maybe one species of toad has just barely what it takes while a nearby species lacks that smidgen extra. Or maybe, on Earth, humans alone have conscious experiences, and the saltation is in evolutionary history, between *Homo habilis*, say, and *Homo erectus*.
2. *Indeterminacy*. Some systems – humans, apes, dogs – are determinately conscious. Other systems – elementary particles, sea sponges – are determinately nonconscious. Somewhere in the middle are systems of intermediate and indeterminate status, capable only of “borderline consciousness” in the sense described above. This indeterminate middle might be relatively low on the scale of neural and cognitive complexity (maybe jellyfish are in the gray area) or it might be relatively high (maybe lizards are in the gray area). This indeterminate middle might be relatively narrow (among currently existing animals, maybe only cnidarians) or it might be broad (maybe all the way from jellyfish to lizards). Either way, no sharp line divides the conscious organisms from the nonconscious ones.
3. *Panpsychism*. We can avoid Horns 1 and 2 by endorsing the view that literally *everything* is conscious. Frogs, snails, lizards, and all the rest – they're all conscious. Dreary monotony might be the sponge's fate, but there's something it's like to be one, waving around in the current, expelling indigestible material. Not only is every animal conscious, but so also is every plant, every computer, every micro-organism, and every elementary particle. (This is not the only form of panpsychism, but it is the form that best avoids Horns 1 and 2 of the quadrilemma.)
4. *Eliminativism*. The other way to avoid Horns 1 and 2 is to deny that consciousness exists at all. “Consciousness”, perhaps, is a broken concept, laden with false presuppositions about, say, immateriality, irreducibility, infallibility, or a subjective “self”. On this way of thinking, the concept is either so beyond repair that we ought to just jettison it or it's a workable enough concept but no entity satisfies the criteria for possessing it.

These four options are logically exclusive and exhaustive. Either nothing is conscious (eliminativism), or everything is conscious (panpsychism), or there's a sharp line between conscious and nonconscious systems (saltation), or there are borderline cases (indeterminacy).

I have framed this quadrilemma in terms of conscious and nonconscious *systems*, since that is more intuitive than some other framings. We can think of a determinately conscious system as any system that sometimes enters determinately conscious states, even if it is not always determinately conscious. A borderline conscious system would then be a system that sometimes enters borderline conscious states but is never determinately conscious.

### 3 Reducing four horns to two

Panpsychism is highly counterintuitive. Most people take for granted that solitary electrons are not conscious. That said, as I've argued elsewhere, something radically contrary to common sense must be true about consciousness (Schwitzgebel, 2014, [forthcoming](#)). Maybe this is it? Furthermore, panpsychism has recently enjoyed a resurgence of interest (e.g., Strawson, 2006; Chalmers, 2015; Goff, 2017; Roelofs, 2019; Seager, 2020).

However, even if we were to accept panpsychism, we might still be forced to choose between saltation and indeterminacy. The theoretical choice can be framed in terms of combinations of fundamental particles. If an electron is a conscious system, is every pair of electrons also a conscious system? If every pair of electrons is a conscious system, is every set of particles in the universe, no matter how arbitrarily scattered, also a conscious system? To say yes is to accept a plentitudinous proliferation of conscious systems in which, for example, every particle in your body participates in maybe  $10^{10^{79}}$  different conscious systems (reflecting every possible combination with the  $10^{80}$  particles in the observable universe). While some panpsychists might be comfortable with this result (Roelofs, 2019, perhaps), most will want to distinguish between combinations of particles that do and do not constitute genuine conscious systems. Only the most radically combinatory panpsychists can avoid the choice between saltation and indeterminacy.<sup>7</sup>

Similarly, eliminativism is highly counterintuitive. And yet, like panpsychism, it has recently enjoyed a resurgence (e.g., Frankish, 2016; Kammerer, 2019), and it might be reasonable to leave space for this view. However, again, even if we were to accept eliminativism, we might still be forced to choose between saltation and indeterminacy. Perhaps the philosophers' technical conception of *phenomenal consciousness* has no proper referent. There just are no immaterial or otherwise spooky properties of the sort supposedly required. Nevertheless, there will be more naturalistically acceptable concepts in the vicinity – perhaps “access consciousness” (Block 1995), or consciousness in Dennett's (1998) sense or the term, or some stripped-down “innocent” conceptualization of consciousness, picked out by example in a way that is adequately neutral about dubious theoretical issues (Schwitzgebel, 2016; Frankish, 2016). Stipulate that *quasi-consciousness* refers to some naturalistically acceptable psychological property that is present in the types of cases that mainstream philosophers regard as paradigmatically conscious and absent in the types of cases that mainstream philosophers regard as paradigmatically nonconscious. If there is more than one such naturalistically acceptable psychological state that meets this criterion, choose the scientifically most useful property. Some eliminativists will be happy enough to use the term “consciousness” to refer to quasi-consciousness, as long as it's clear that what is meant is not “phenomenal consciousness” in an objectionably laden sense. It follows from the stipulative definition that some systems will

<sup>7</sup> Roelofs proposes the term “intelligent subject” (2019, p. 151) for a subclass of conscious subjects – approximately the class of entities we intuitively regard as conscious. Roelofs acknowledges that “intelligent subject” is a vague term (p. 156). If it can be argued that the term “conscious” does or should mean “intelligent subject”, then Roelofs' combinatorial panpsychism might be open to a redefinition maneuver analogous to the one I apply to eliminativism in the next paragraph.

be quasi-conscious in this sense and some will not. Thus, the eliminativist escapes the choice between saltation or indeterminacy for the concept whose usefulness they reject (“phenomenal consciousness”) but faces the choice for a nearby concept whose usefulness they presumably would accept (“quasi-consciousness” or “consciousness” in an innocent sense). For the sake of the present argument, I ask them to interpret me as employing the latter concept when I use the term “conscious”.

Perhaps either literally every mereological sum of scattered things is conscious, or absolutely nothing is. But if we reject both extreme views, there must be a division somewhere, between systems or states that are conscious and those that are not. The division is either sharp, like that between being in the ground state orbital or not, or it is vague, like that between green and not green. Saltation or indeterminacy.

#### **4 Contra saltation, part one: consciousness is a categorical property with (probably) a graded basis**

Baldness is a categorical property with a graded basis. A person is either determinately bald, determinately non-bald, or in the gray area between. In that sense, baldness is categorical. However, the basis or grounds of baldness is graded: number of hairs and maybe how long, thick, and robust those hairs are. If you have enough hair, you’re not bald, but there’s no single best place to draw the categorical line. Similarly, greenness and extraversion are categorical properties with graded bases that defy sharp-edged division. In contrast, being in the ground orbital is a categorical property without a graded basis. That’s the “quantum” insight in quantum theory. Bracketing cases of superposition, the electron is either in this orbital, or that one, or that other one, discretely. There’s discontinuity as it jumps, rather than gradations of close enough. Similarly, although the real numbers are continuous, a three followed by any finite number of nines is discretely different from exactly four. Being *approximately* four has a graded basis, but being *exactly* four is sharp-edged.

Most naturalistic theories of consciousness give consciousness a graded basis. Consider broadcast theories, like Dennett’s “fame in the brain” theory (2005; similarly Tye, 2000; Prinz, 2012). On such views, a cognitive state is conscious if it is sufficiently “famous” in the brain – that is, if its outputs are sufficiently well-known or available to other cognitive processes, such as working memory, speech production, or long-term planning. Fame, of course, admits of degrees. *How much* fame is necessary for consciousness? And in what respects, to what systems, for what duration? There’s no theoretical support for positing a sharp, categorical line such that consciousness is determinately absent until there is exactly this much fame in exactly these systems (see Dennett, 1998, p. 349; Tye, 2000, p. 180–181).

Global Workspace Theories (Baars, 1988; Dehaene, 2014; Mashour et al., 2020) similarly treat consciousness as a matter of information sharing and availability across the brain. This also appears to be a matter of degree. Maybe typically once a process crosses a certain threshold it tends to quickly become widely available in a manner suggestive of a phase transition. Nevertheless, measured responses and brain activity are sometimes intermediate between standard “conscious” and “nonconscious”

patterns.<sup>8</sup> Looking at non-human cases, the graded nature of Global Workspace theories is even clearer (Carruthers, 2019). Even entities as neurally decentralized as jellyfish and snails employ neural signals to coordinate whole-body motions. Is that “workspace” enough for consciousness? Artificial systems, also, could presumably be designed with various degrees of centralization and information sharing among their subsystems. Again, there’s no reason to expect a bright line.

Or consider a very different class of theories, which treat animals as conscious if they have the right kinds of general cognitive capacities, such as “universal associative learning”, trace conditioning, or the ability to match opportunities with needs using a central motion-stabilized body-world interface organized around a sensorimotor ego-center (Merker, 2007; Ginsburg and Jablonka 2019; Birch, 2022). These too are capacities that come in degrees. How flexible, exactly, must the learning systems be? How long must a memory trace be capable of enduring in a conditioning task, in what modalities, under what conditions? How stable must the body-world interface be and how effective in helping match opportunities with needs? Once again, the categorical property of conscious versus nonconscious rests atop what appears to be a smooth gradation of degrees, varying both within and between species, as well as in evolutionary history and individual development.

Similarly, “higher-order” cognitive processes, self-representation, attention, recurrent feedback networks, even just having something worth calling a “brain” – all of these candidate grounds of consciousness are either graded properties or are categorical properties (like having a brain) that are in turn grounded in graded properties with borderline cases. Different species have these properties to different degrees, as do different individuals within species, as do different stages of individuals during development. Look from one naturalistic theory to the next – each grounds consciousness in something graded. On such views, an entity is conscious if it has *enough* of property X, where X depends on which theory is correct, and where “enough” is plausibly a vague matter. There are few truly sharp borders in nature.

Apart from rejecting naturalistic theories of consciousness of this sort, I see two ways to resist this conclusion, which I will call the *Perfectly Sharp Phase Transition View* and the *Luminous Penny View*.

## 5 Contra saltation, part two: against the Perfectly Sharp Phase Transition View

Water cools and cools, not changing much, then suddenly it solidifies into ice. The fatigued wooden beam takes more and more weight, bending just a bit more with each kilogram, then suddenly it snaps and drops its load. On the Perfectly Sharp Phase Transition View, consciousness is like that. The basis of consciousness might admit of degrees, but still there’s a sharp and sudden transition between nonconscious and conscious states. When water is at 0.1° C, it’s just ordinary liquid water. At 0.0°, something very different happens. When the Global Workspace (say) is size X-1, sure, there’s a functional workspace where information is shared among subsystems,

<sup>8</sup> Sergent et al., 2005; Salti et al., 2015; Sy et al., 2021.

there's unified behavior of a sort, but no consciousness. When it hits X – when there's that one last crucial neural connection, perhaps – bam! Suddenly everything is different. The bright line has been crossed. There's a phase transition. The water freezes, the beam snaps, consciousness illuminates the mind.

I'll present a caveat, a dilemma, and a clarification.

The caveat is: Of course the water doesn't *instantly* become ice. The rod doesn't *instantly* snap. If you zoom in close enough, there will be intermediate states. The same is likely true for the bases of consciousness on naturalistic views of the sort discussed above, unless those bases rely on genuine quantum-level discontinuities. Someone committed to the impossibility of borderline cases of consciousness even in principle, even for an instant, as a matter of strict necessity, ought to pause here. If the phase transition from nonconscious to conscious needs to be *truly* instantaneous without a millisecond of in-betweenness, then it cannot align neatly with any ordinary, non-quantum, functional or neurophysiological basis. It will need, somehow, to be sharper-bordered than the natural properties that ground it.

The dilemma is: The Perfectly Sharp Phase Transition View is either empirically unwarranted or it renders consciousness virtually epiphenomenal.

When water becomes ice, not only does it change from liquid to solid, but many of its other properties change. You can cut a block out of it. You can rest a nickel on it. You can bruise your toe when you drop it. When a wooden beam breaks, it emits a loud crack, the load crashes down, and you can now wiggle one end of the beam without wiggling the other. Phase transitions like this are notable because many properties change suddenly and in synchrony. But this does not appear always to happen with consciousness. That precipitates the dilemma.

There are phase transitions in the human brain, of course. One is the transition from sleeping to waking. Much changes quickly when you awaken. You open your eyes and gather more detail from the environment. Your EEG patterns change. You lay down long-term memories better. You start to recall plans from the previous day. However, this phase transition is *not* the phase transition between nonconscious and conscious, or at least not as a general matter, since you often have experiences in your sleep. Although people sometimes say they are “unconscious” when they are dreaming, that's not the sense of consciousness at issue here, since dreaming is an experiential state. There's something it's like to dream. Perhaps there is a phase transition between REM sleep, associated with longer, narratively complex dreams, and NREM sleep. But that probably isn't the division between conscious and nonconscious either, since people often also report dream experiences during NREM sleep. Similarly, the difference between being under general anesthesia and being in an ordinary waking state doesn't appear to map neatly onto a sharp conscious/nonconscious distinction, since people can apparently sometimes be conscious under general anesthesia and there appear to be a variety of intermediate states and dissociable networks that don't change instantly and in synchrony, even if there are also often rapid phase transitions.<sup>9</sup>

While one could *speculate* that all of the subphases and substates of sleep and anesthesia divide sharply into determinately conscious and determinately noncon-

<sup>9</sup> John and Pritchep, 2005; Bruno et al., 2011; Lee et al., 2011; Casali et al., 2013; Bonhomme et al., 2019.



scious, the empirical evidence does not provide positive support for such a view. The Perfectly Sharp Phase Transition View, to the extent it models itself on water freezing and beams breaking, is thus empirically unsupported in the human case. Sometimes there are sudden phase transitions in the brain. However, the balance of evidence does not suggest that falling asleep or waking, starting to dream or ceasing to dream, falling into anesthesia or rising out of it, is always a sharp transition between conscious and nonconscious, where a wide range of cognitive and neurophysiological properties change suddenly and in synchrony. The Perfectly Sharp Phase Transition View, if intended as a defense of saltation, is committed to a negative existential generalization: There can be *no* borderline cases of consciousness. This is a very strong claim, which fits at best uneasily with the empirical data.

Let me emphasize that last point, by way of clarification. The Perfectly Sharp Phase Transition View as articulated here commits to the nonexistence of borderline consciousness *ever*. That is much bolder than any empirical claim that transitions from nonconscious to conscious states are *typically* phase-like. My argument here in no way conflicts with empirical claims by, for example, Lee et al. (2011) and Dehaene (2014) that sudden phase transitions are typical and important.

The Perfectly Sharp Phase Transition View looks empirically even weaker when we consider human development and non-human animals. It *could* have been the case that when we look across the animal kingdom we see something like a “phase transition” between animals with and without consciousness. These animals over here have the markers of consciousness and a wide range of corresponding capacities, and those animals over there do not, with no animals in the middle. Instead, non-human animals have approximately a continuum of capacities. Similarly, in human development we *could* have seen evidence for a moment when the lights turn on, so to speak, in the fetus or the infant: Consciousness arrives, and instantly everything is visibly different, behaviorally and neurophysiologically. But there is no evidence of such a saltation. Even birth is a temporally extended process.

That's the first horn of the dilemma for the Perfectly Sharp Phase Transition View: Accept that the sharp transition between nonconscious and conscious should be accompanied by the dramatic and sudden change of many other properties, then face the empirical evidence that the conscious/nonconscious border does not always involve a sharp, synchronous, wide-ranging transition. The Perfectly Sharp Phase Transition View can escape by retreating to the second horn of the dilemma. According to this second horn, consciousness is cognitively, behaviorally, and neurophysiologically unimportant. Although *consciousness* always transitions sharply and dramatically, nothing else need change much. The lights turn on, but the brain need hardly alter. The lights turn on, but there need be no correspondingly dramatic shift in memory, or attention, or self-knowledge, or action planning, or sensory integration, or... All of the latter can still change slowly or asynchronously, in accord with the empirical evidence.

This view is unattractive for at least three reasons. First, it dissociates consciousness from its naturalistic bases. We began by thinking that consciousness is information sharing or self-representation or whatever, but now we are committed to saying that consciousness can change radically in a near-instant, while information sharing or self-representation or whatever hardly changes at all. Second, it dissociates con-

consciousness from the evidence for consciousness. The evidence for consciousness is, presumably, performance on introspective or other cognitive tasks, or neurophysiological conditions associated with introspective reports and cognitive performance; but now we are postulating big changes in consciousness that elude such methods. Third, most readers, I assume, think that consciousness is important, not just intrinsically but also for its effects on what you do and how you think. But now consciousness seems not to matter much.<sup>10</sup>

The Perfectly Sharp Phase Transition View postulates a qualitative change between nonconsciousness and consciousness, like the change from liquid to solid, except perfectly sharp, with no borderline cases. It's this big change that precipitates the dilemma, since either the Perfectly Sharp Phase Transition advocate should also expect there always also to be sudden, synchronous cognitive and neurophysiological changes (in conflict with the most natural reading of the empirical evidence) or they should not expect such changes (making consciousness approximately epiphenomenal).

The saltationist can attempt to escape these objections by denying that the sharp border involves a big change in consciousness. It might instead involve the discrete appearance of a tiny smidgen of consciousness. This is the Luminous Penny View.

## 6 Contra saltation, part three: against the Luminous Penny View

Being conscious might be like having money. You might have a little money, or you might have a lot of money, but having any money at all is discretely different from having not a single cent.<sup>11</sup> Maybe a sea anemone has just a tiny bit of consciousness, a wee flicker of experience – at one moment a barely felt impulse to withdraw from something noxious, at another a general sensation of the current sweeping from right to left. Maybe that's \$1.50 of consciousness. You, in contrast, might be a consciousness millionaire, with richly detailed consciousness in several modalities at once. However, both you and the anemone, on this view, are discretely different from an electron or a stone, entirely devoid of consciousness.<sup>12</sup>

Imagine the visual field slowly collapsing. It shrinks and shrinks until nothing remains but a tiny gray dot in the center. Finally, the dot winks out. In this way,

<sup>10</sup> Could we arrange atop a graded basis, such as workspace size, a series of properties, consciousness<sub>1</sub>, consciousness<sub>2</sub>, consciousness<sub>3</sub>, etc., exactly one of which is the real property of consciousness? We cannot escape the dilemma this way. Either something crucially important changes with the tiny increase in workspace size, or nothing crucially important changes. The first option throws us back against the empirical evidence described in the body of the article. The second option leaves us no reason to care specifically about the property of consciousness as opposed to other nearby properties and probably requires commitment to a theory of vagueness on which this is exactly the type of thing we should expect for borderline cases of vague properties in general – thus conforming to rather than conflicting with my overall thesis (see note 5).

<sup>11</sup> There are probably borderline cases of money-possession too, such as when your last dollar is in transit to a creditor or when all your money is in the currency of a collapsing regime, but set such cases aside for the sake of the analogy.

<sup>12</sup> See Lee, [forthcoming](#), for a general discussion of degrees of consciousness that is sensitive to the distinction made in this paragraph.

there might be a quantitative difference between lots of visual consciousness and a minimum of it, and then a discontinuous qualitative difference between the minimum possible visual experience and none at all.

On the *Luminous Penny View*, as I'll call it, there is a saltation from nonconscious to conscious in the sense that there are no in-between states in which consciousness is neither determinately present nor determinately absent. Yet the saltation is to such an impoverished state of consciousness that it is almost empirically indistinguishable from lacking consciousness. Analogously, in purchasing power, having a single penny is almost empirically indistinguishable from complete bankruptcy. Still, that pennysworth of consciousness is the difference between the "lights being on", so to speak, and the lights being off. It is a luminous penny.<sup>13</sup>

The view escapes the empirical concerns that face the Phase Transition View, since we ought no longer expect big empirical consequences from the sudden transition from nonconscious to conscious. However, the Luminous Penny View faces a challenge in locating the lower bound of consciousness, for both states and animals.

Start with animals. What kind of animal would have only a pennysworth of consciousness? A lizard, maybe? That seems an odd view. Lizards have complex visual capacities. If they are visually conscious at all, it seems natural to suppose that their visual consciousness would approximately match their visual capacities – or at least that there would be *some* visual complexity, more than the minimum possible, more than a tiny gray dot. It's equally odd to suppose that a lizard would be conscious without having visual consciousness. What would its experience be? A bare minimal striving, even simpler than the states imaginatively attributed the anemone a few paragraphs back? A mere thought of "here, now"?

More natural is to suppose that if a lizard is determinately conscious, it has more than the most minimal speck of consciousness. To find the minimal case, we must then look toward simpler organisms. How about ants? Snails? The argument repeats: These entities have more than minimal sensory capacities, so if they are conscious it's reasonable to suppose that they have sensory experience with some detail, more than a pennysworth. Reasoning of this sort leads Chalmers (1996) to a panpsychist conclusion: The simplest possible consciousness requires the simplest possible sensory system, such as the simple too-cold/okay of a thermostat.

The Luminous Penny View thus faces its own dilemma: Either slide far down the scale of complexity to a position nearly panpsychist (though perhaps not as radically panpsychist as the every-combination view) or postulate the existence of some middle-complexity organism that possesses a single dot of minimal consciousness despite having a wealth of sensory sensitivity.

Perhaps the problem is in the initial move of quantifying consciousness, that is, in the commitment to saying that complex experiences somehow involve "more" consciousness than simple experiences? Maybe! But if you drop that assumption, you drop the luminous penny solution to the problem of saltation.

---

<sup>13</sup> One might also imagine a Luminous Nickel View or a Luminous \$1.50 View, with sudden saltation to a somewhat impoverished but not minimal state. This is a compromise between the Luminous Penny View and the Perfectly Sharp Phase Transition View, inheriting the difficulties of both, though perhaps in somewhat reduced form. (Thanks to Henry Shevlin for this suggestion.)

State transitions in adult humans raise a related worry. We have plausibly nonconscious states on one side (perhaps dreamless sleep<sup>14</sup>), indisputably conscious states on the other side (normal waking states), and complex transitional states between them that lack the kind of simple structure one might expect to produce exactly a determinate pennynorth of consciousness and no more.

If consciousness requires sophisticated self-representational capacity (as, for example, on “higher order” views), lizard or garden snail consciousness is presumably out of the question. But what kind of animal, in what kind of state, would have exactly one self-representation of maximally simple content? (Only always “I exist” and nothing more?) Self-representational views fit better either with phase transition views or with gradualist views that allow for periods of indeterminacy as self-representational capacities slowly take shape and, to quote Wittgenstein, “light dawns gradually over the whole” (Wittgenstein 1951/1969, § 141).<sup>15</sup>

If you’re looking for a penny, ask a panpsychist (or a near cousin of a panpsychist, such as an Integrated Information Theorist: Oizumi et al., 2014). Maximally simple systems are the appropriate hunting grounds for maximally simple consciousness, if such a thing as maximally simple consciousness exists at all. From something as large, complicated, and fuzzy-bordered as brain processes, we ought to expect either large, sudden phase transitions or the gradual fade-in of something much richer than a penny.

## 7 On the inconceivability of borderline consciousness

Borderline consciousness might seem inconceivable. We can imagine that there’s something it’s like to be a particular garden snail at a particular moment, or we can imagine that there’s nothing it’s like, but it seems impossible to imagine its *sort of* being like something. How might such an in-between state feel, for the snail? As soon as we try to answer that question, we seem forced either to say that it wouldn’t feel like anything or to contemplate various types of conscious experiences the snail might have. We can imagine the snail’s having some flow of experience, however limited, or we can imagine the snail to be an experiential blank, but we can’t in the same way imagine some in-between state such that it’s neither determinately the case that the snail has conscious experiences nor determinately the case that the snail lacks conscious experiences. Even a dim light is a light, determinately on, and even a flickering light is determinately on or off at any particular moment.

Similarly, we can imagine what it’s like to be in a hazy, confused state immediately after waking. We can imagine, and maybe remember, feeling unsure where we are, being unable to clearly distinguish dream from reality, and not yet having a sense of the coming day. Maybe that’s what it’s like, sometimes, to slowly awaken. But in imagining this, we seem to be imagining a state that is determinately conscious (though it might have some indeterminate content in the sense of Sect. 1). We

<sup>14</sup> Though see Windt et al., 2016, for discussion of non-dream experiences during sleep.

<sup>15</sup> On gradualism in cognitive development, see also Schwitzgebel, 1999; McGeer and Schwitzgebel, 2006.

can also, in a different way, imagine being in a dreamless sleep with no experiences whatsoever – though of course we can't imagine what it's like to be in such a deep sleep, because there's nothing it's like. As soon as we try to imagine the transition between nonconscious sleep and dream, or between nonconscious sleep and waking, we start to imagine dream experiences or waking experiences or confused half-awake experiences, that is, experiences of some sort or other. We imagine that it's like nothing – nothing – nothing – something – something – something. Between nothing and something is no middle ground of half-something. A half-something is already a something. Borderline consciousness, it seems, must already be consciousness unless it is no consciousness at all.

I grant that all of this is highly intuitive. Such considerations, on their face, present a serious obstacle to understanding what could be meant by “borderline consciousness”. What kind of state could the phrase even refer to? Before admitting the existence of borderline or in-between states of consciousness, we want to know what such a state *would be like*. We want a sense of it, a feel for it. We want to remember some borderline experiences of our own. Before accepting that a snail or frog might be borderline conscious, neither determinately lights on nor determinately lights off, we want at least a speculative gesture toward the experiential character of such in-betweenish phenomenology.

I feel the pull of this manner of thinking, but it is a paradoxical demand. It's like the Catch-22 of needing to complete a form to prove that you're incompetent, the completing of which proves that you're competent. It's like demanding that the borderline shade of only-kind-of-green must match some sample of determinate green before you're willing to accept that it's a borderline shade that doesn't match any such sample. An implicit standard of conceivability drives the demand, which it is impossible to meet without self-contradiction. The implicit standard appears to be this: Before granting the existence of borderline consciousness, we want to be able to imagine what it would be like to be in such a state. But of course there is not determinately anything it's like to be in such a state. The more we try to imagine what it would be like, the worse we miss our target. If you look through a filter that shows only determinately bald people, you won't see anyone who is borderline bald. But you ought not conclude that no borderline bald people exist. The fault is in the filter. The fault is in the imaginative demand.

In another sense, borderline states of consciousness are perfectly conceivable. They're not like four-sided triangles. There is no self-contradiction in the very idea. If you're unhappy with your understanding of them, it could be because you desire something that you cannot reasonably expect to have. The proper response is to shed the desire.

A philosophically inclined middle-schooler, on their first introduction to imaginary numbers, might complain that they cannot conceive of a number whose square is  $-1$ . What is this strange thing? It fits nowhere on the number line. You can't hold  $3i$  pebbles. You can't count  $3i$  sheep. So called “imaginary numbers” might seem to this middle-schooler to be only an empty game with no proper reference. And yet there is no contradiction in the mathematics. We can use imaginary numbers. We can even frame physical laws in terms of them, as in quantum mechanics. In a *certain way*, imaginary numbers are, despite their name, unimaginable. But the implicit criterion

of imagination at work – picturing 3i sheep, for example – is inappropriate to the case.

Usually it is unreasonable to demand something incoherent. But one exception to this rule is in a *reductio*. If there were good reason to think that all numbers *must* be imaginable or conceivable on the real number line, then the fact that imaginary numbers have no place on the real number line would be a mark against them. The middle-schooler's imaginative task would still be incoherent, but the incoherence would be, so to speak, the fault of a flaw in the concept of imaginary numbers. This opens an argumentative avenue for those who would deny the existence of borderline consciousness. If something in the nature of borderline consciousness requires that we be able to succeed in the incoherent task of imagining or remembering determinately what it is like to be in a state that it is not determinately like anything to be in, then it would make sense on these grounds to reject borderline consciousness. Sect. 9 will discuss some arguments in this direction.

Admittedly, I haven't provided much by way of a positive conception of borderline states of consciousness, apart from gesturing toward clear cases of consciousness and clear cases of nonconsciousness and saying that there is good reason to accept a gray area between. If the situation were different in one of two ways, I could do more.

First, if I had the right theory of consciousness on hand, I could characterize borderline cases in terms of that theory. Suppose, for example, that I knew that Global Workspace Theory was true, and suppose I knew that when information occupies a workspace of size X it's determinately conscious and when information occupies a workspace of size X-1000 it's determinately nonconscious. Suppose also that the theory implies no sudden saltation between. I could then point at workspaces of size X-360 and X-742, theorize about them, discuss how they work in animal cases and maybe some human state transitions, and in that way render the idea familiar and workable. Someone initially skeptical might slowly lose the feeling that they must be able to imagine or remember what borderline conscious states are like to accept their existence.

Second, if borderline consciousness were common enough and important enough in human life, we might grow accustomed to the idea and even develop an ordinary language term for it. We might say, "ah yes, one of those *mizzy* states, in the intermediate zone between consciousness and nonconsciousness." But we have no need for such a folk concept. We care little about and needn't track borderline cases, whether in transition between dreamless sleep and dreams, dreamless sleep and waking, or in and out of anesthesia or other forms of nonconsciousness. For everyday purposes, we can be loose in our thinking about the twilight zone. The same goes for development and phylogeny. We assume that infants are determinately conscious from birth, so any developmental borderline consciousness is fetal. Nonhuman animals we can adequately enough imagine as either determinately conscious or as nonconscious machines. There has never been serious linguistic pressure toward accurate phenomenological representation of fetuses and nonhuman animals.

We will thus need to reconcile ourselves with a certain sort of dissatisfaction. It's incoherent to attempt to imagine, in any determinate way, what it would be like to be in a borderline conscious state, since there's nothing determinate it would be like; so first-person imaginative transportation and phenomenological memory won't pro-

vide a good handle on the idea. But neither can we yet gain a good handle through third-person methods in consciousness science, since consciousness science is still so contentious and underdeveloped. Nor do we have a pre-existing folk concept or any habit of thinking carefully about such states. All we can do is indicate uncontroversially determinate cases of consciousness over here and uncontroversially determinate cases of nonconsciousness over there and say, somewhere in the middle, borderline cases must exist.

## 8 Two types of indeterminacy about consciousness

Papineau (2002, 2003) has argued at length for indeterminacy not only concerning the contents of consciousness but also concerning “consciousness as such”. In this respect, his view is similar to mine. Consciousness, Papineau argues, can be identified with any one of a number of material properties, such as a certain neurophysiological state, a certain functional state that always co-occurs with the neurophysiological state in ordinary cases, a certain first-order representational state, or a certain higher-order representational state. The term “consciousness” is insufficiently precisely specified to pick out exactly one of those properties as opposed to another (similarly Chin, 2015; Birch, 2021). This type of indeterminacy is different from the indeterminacy of in-betweenness. Citing Block (2002/2007), Papineau writes:

Actually, an analogy with baldness is less than perfect, since baldness is a matter of degree, vague because there is no sharp cut-off point. Phenomenal concepts, by contrast, are vague because it is indeterminate what kind of property they refer to, rather than where to divide some continuum. A better analogy is provided by the traditional [Inuit] word for whale oil.... At some point in [Inuit] history, a petroleum product which looked and functioned just like natural whale oil was introduced as a substitute. Did the original [Inuit] word (let's pretend it was 'whale oil') apply to this new substance or not? If 'whale oil' referred to a biologically or chemically identified type, it did not; but if 'whale oil' referred to anything with the requisite appearance and use, then it did. But of course there is likely to be no answer here (2003, p. 215).

Let's call the whale-oil species of indeterminacy *kind indeterminacy* and indeterminacy due to vagueness along a continuum *in-betweenness indeterminacy*.<sup>16</sup> Other cases of kind indeterminacy might include whether H<sub>2</sub>O composed of heavy hydrogen (that is, D<sub>2</sub>O) is water or not, or whether tree-like organisms on an alien planet would be trees or not. Papineau defends only kind indeterminacy for consciousness

<sup>16</sup> In-betweenness indeterminacy in the present sense includes both degree and combinatory indeterminacy in Alston's (1964) sense. It might also be indeterminate whether a particular kind of indeterminacy is kind indeterminacy or in-betweenness indeterminacy, for example, whether some particular song is “rock music” or not. I take no stand on whether kind indeterminacy involves vagueness or not (e.g., Weatherson, 2009).

and might in fact reject in-betweenness indeterminacy, as suggested by his contrasting consciousness with baldness in the quote above.<sup>17</sup>

Although Papineau and I agree that “consciousness” is a term that admits indeterminate cases, my view is the opposite of Papineau’s regarding the type of indeterminacy. My arguments above concern only *in-betweenness indeterminacy*. The existence of *kind indeterminacy* is less clear to me. There is, as I suggested in Sect. 1, an obvious property that conscious experiences share in common, and the term “conscious” refers to that obvious property. I see little room for its being indeterminate what kind of property this term refers to – though of course we might be, and probably are, ignorant about what physical or functional property is identical with this introspectively discoverable property, or whether indeed any physical or functional property is identical to it. The kind in question – consciousness – is, I hope, a determinate or close to determinate kind. I argue for indeterminacy only regarding the boundaries or borders of that determinate kind. However, if it turns out that consciousness *also* shows kind indeterminacy, that does not impair my main argument, as long as each candidate kind itself admits of borderline cases, as would seem to be plausible if each has a graded basis.

## 9 Antony and Simon against vagueness

Although many philosophers have expressed reservations about the possibility of borderline cases of consciousness (see note 2), I am aware of only two sustained arguments against it, the arguments of Antony (2006, 2008) and Simon (2017). Both are versions of the inconceivability objection.

According to Antony, to be competent with a vague concept F, we must be able to conceive of clear Fs, clear not-Fs, and borderline Fs. Conceiving of borderline Fs requires being able to represent a series of individual cases from not-F to F in which the individuals’ parts or properties gradually change from being such that the individual is clearly not-F to being clearly F with borderline cases between. For example, being competent with the concept *bald* requires being able to imagine a series of people with varying amounts of hair from clearly not-bald to clearly bald with borderline bald people between. Antony argues that this condition cannot be met for the concept *conscious state*. The reason is that there are no parts or properties that can be arranged in such a series. There is no analog of the amount of hair. Some candidate elements, like having qualia or a point of view, entail the presence of consciousness and so cannot be arranged in such a series. Other candidate elements, such as having physical or functional structure such-and-such, are not theoretically neutral in the way required for the concept of consciousness, since our current concept of consciousness is theoretically neutral among various materialist, dualist, and idealist theories. Consequently, we cannot conceive of borderline cases in the required way.

<sup>17</sup> Although I don’t find the phrasing completely clear, Papineau also says, regarding a “silicon doppelgänger” (conscious on a functional-structural theory but not on a neurophysiological theory) that although it’s a vague case of consciousness “It is not vague how it is for the doppelgänger himself” (2003, p. 219; cf. 2002, p. 227). In earlier work, Papineau endorses in-betweenness indeterminacy along the way to defending kind indeterminacy (1993, p. 124).



Antony considers a hypothetical objector who argues that borderline conscious states can't be imagined because they can't be introspected and thus that our failure to clearly conceive of borderline cases is "an artifact of the nature of introspection, which has nothing to do with whether there are borderline cases" (2008, p. 260). This is close to the objection I would make, drawing on considerations explored in Sect. 7. The demand that we clearly conceive or imagine borderline cases of consciousness is illegitimate, if the standard of adequate conception or imagination requires conceiving or imagining determinately *what it would be like* to be in a borderline state. By definition, such states are not determinately like anything. To this hypothetical objector, Antony replies that failure of introspection does not show that the boundary is in fact vague. He then reiterates the claim that imagining the boundary between the conscious and the nonconscious is a matter of imagining what such states "are like", which yields an intuitively sharp boundary between the presence and absence of phenomenology.

Two initial remarks: First, I am happy to grant that the impossibility of clearly introspecting borderline consciousness doesn't show that the conscious-nonconscious boundary is vague. My case rests on entirely different grounds. Second, it is telling that Antony's reply is *not* to change the criterion of conceivability to something other than imagining what the states would be phenomenologically like, as he might do if he were aiming to avoid the illegitimately paradoxical demand. Instead, Antony appears to repeat the illegitimate demand.

However, let's grant that the demand is in a certain sense legitimate. Maybe, to be fully competent with a vague concept, we need a good sense of the borderline cases, and maybe for the concept of consciousness we don't have a good enough sense. It then follows that *either* the concept of consciousness is not vague *or* we are not fully competent with that concept. We may choose the second disjunct. We have only a partial conception of consciousness, based on our ability to conceive of clear cases of its presence or absence, but limited by our incomplete grasp of the borderline cases that must (if the arguments of previous sections are correct) exist. A concept with which we aren't fully competent by Antony's high standards might nonetheless successfully pick out an interesting property in nature. In the year 1500, the human concepts of *gold*, *planet*, and *alive* left much to be desired. Those concepts were grounded in examples, with a poor understanding of the underlying nature and range of possible cases, and they were in some respects erroneous. Yet the concepts functioned well enough, successfully referring to entities in nature. People might have intuitively regarded *gold* as a graded concept, admitting a range of borderline cases from lead to gold, and they might intuitively have regarded *alive* as a sharply-bordered concept, but the phenomena themselves turned out to be closer to the reverse.<sup>18</sup> At best, Antony shows that we don't fully grasp *conscious* as a vague-boundaried concept – a fact that might concern only our current limitations – not that

<sup>18</sup> Continuing the parallel, one might imagine an analog of the Perfectly Sharp Phase Transition View and Luminous Penny View for a saltationist conception of "alive", leading to similar objections. Although there is a phase transition between alive and not alive, the evidence suggests that it is not always perfectly sharp, unless it is nearly epiphenomenal. Similarly, a Luminous Penny analog for alive might generate problems in locating the lower bound. Despite such difficulties, reference magnetism (or something similar) ensured that the concept picked out a real natural property.

consciousness, as it actually exists in nature, is a sharp-boundaried phenomenon. Indeed, drawing a historical parallel with the concept of life, Antony suggests that in the future we might shift from having a sharp to a graded concept as our empirical knowledge improves. Thus, Antony and I needn't fundamentally disagree.

Simon's (2017) argument that *phenomenally conscious* has no borderline cases also turns on conceivability. Simon holds, similarly to Antony, that for the predicate "phenomenally conscious" to be vague, there must be a "positive characterization" of a borderline case that can show any sufficiently competent speaker what makes it a borderline case. Simon argues that no purely physical or structural description could suffice, since no such description could render it knowable to any competent speaker that the case is a borderline case of "phenomenally conscious".<sup>19</sup> For an adequate characterization, then, we must turn to cases like slowly falling asleep, or cases of consciousness without attention, or cases of hearing ever quieter sounds until one is no longer sure whether one is hearing a sound or not (Tye, 1996). However, Simon suggests that no such case is convincing as a positive characterization of an in-between state of borderline consciousness, since each description is consistent with there being determinate cutoffs.

The primary weakness of the argument, if construed as an objection to my thesis, is the same as Antony's. It moves too quickly from our inability to adequately characterize or conceptualize borderline cases to their probable nonexistence. At best, these arguments reveal limits in our *conceptualization* of consciousness.

Antony or Simon might object as follows. For many vague properties, such as baldness, we can conceive of two clear cases of the property that differ in such a way that one is closer to being a borderline case than the other. For example, we might imagine two determinately bald men, one of whom has slightly more hair than the other and thus is closer to being a borderline case. But it's not clear we can do the same with consciousness. Perhaps no two determinately conscious states will be such that one is closer to being nonconscious than the other; they will both be just fully conscious, full stop.<sup>20</sup>

To this concern, I see two response options. The first option is as follows. Some vague-boundaried properties consist mostly of clear cases intuitively equidistant from the borderline. The many cars, motorcycles, bicycles, trucks, scooters, and buses on our roads are (plausibly) all just straight-up cases of *vehicle*, with none closer to the borderline than the other, even though there are borderline cases of vehicle (a motor home with its wheels off, a piece of cardboard being used as a sled). Similarly, humans, worms, bacteria, and trees are all (plausibly) equally *living things*, with no one type closer to being a borderline case of a living thing than the other. If that's correct for at least some vague-boundaried properties, then consciousness might be like that, with the intuitive, familiar cases all equally far from borderline. For conscious-

<sup>19</sup> Simon thus relies on assuming the conceivability of "zombies" that are physically and structurally identical to human beings but who lack consciousness. I will not challenge that assumption here. Like Antony, Simon (p. 2121) imagines the possibility of people adopting new concepts that permit positive characterizations of borderline cases.

<sup>20</sup> Thanks especially to Ryan McElhaneey for helpful conversation on this point.

ness, the borderline cases (unlike for vehicle or living thing) might then be hard to locate and conceptualize for the reasons discussed in Sect. 7.

The second option is to take a cue from the Luminous Penny View. Even though there might be no luminous pennies, it might still be reasonable to say that determinately conscious states typically differ in *how conscious* they are. Perhaps one is more vivid than another. Perhaps one is more readily introspectible than another. Perhaps one is more widely broadcast or accessible than another. The less vivid state, less introspectible state, the less widely broadcast or accessible state, or the state with less of some other Feature X might thus be closer to being a borderline case, if borderline states of consciousness are states with insufficient vividness, introspectibility, broadcast, accessibility, or Feature X. If the property in question is experiential or phenomenal, maybe this is a path to partial positive phenomenological conceivability of borderline consciousness – though I prefer not to hang my argument on this idea.

If we had no good grounds for thinking that consciousness admits borderline cases and no good explanation for our failure to be able to clearly characterize borderline cases, then the lack of a clear characterization might justify doubt about the existence of such cases. However, there are good reasons to think that consciousness does admit borderline cases, namely, (1) the graded nature of the naturalistic phenomena that presumably constitute or give rise to conscious experiences and (2) the problems that confront the other three horns of the quadrilemma, that is, panpsychism, eliminativism, and saltation. And our current inability to clearly characterize such states can be explained. It can be explained by, on the one hand, our temptation to paradoxically want to know what they are like, which invites but also ruins the imaginative or memorial approach, and, on the other hand, by the lack of a good folk or scientific understanding of consciousness in general and borderline conscious states in particular.<sup>21</sup>

Is it really possible that our conceptualization of consciousness is as imperfect as this response to Antony and Simon requires? Goff (2017), for example, argues that our knowledge of consciousness is so excellent that the essence of any conscious state is revealed upon attending to that state and forming a “direct phenomenal concept” of it (a concept based wholly on attending to that state). If this were so, then perhaps there would be no room for ignorance of borderline cases. However, most mainstream naturalists about consciousness would presumably reject the strong form of “revelation” required. Introspection is imperfect (e.g., Schwitzgebel, 2011) and fails to reveal the material or functional naturalistic basis of the underlying phenomena (e.g., Smart, 1959). Nor is it clear, in general, that grasping essences requires knowing about borderline cases.<sup>22</sup> Maybe we have excellent knowledge of experiences we introspectively attend in the right way. But borderline cases might be rare,

<sup>21</sup> See Hall [forthcoming](#) for a very different critique of Antony and Simon, based primarily on the logic of vagueness and the relation between knowledge claims and vagueness.

<sup>22</sup> Goff, 2013 argues that if we accept both the transparency thesis associated with property dualism and a linguistic theory of vagueness, then all the sharpenings of the term “consciousness” ought to be a priori knowable. I’m inclined to think that only on an extremely strong version of transparency would a linguistic theory of vagueness deliver the a priori accessibility of all sharpenings. Compare with knowing that the essence of gold is to have 79 protons (if that is indeed its essence). We might know this, be quite competent with the concept of gold, and yet not have a clear conceptualization of how to sharpen borderline cases or

nonhuman, or introspectively elusive – just the types of states that ordinary humans in ordinary conditions *won't* ordinarily discover when they turn their attention inward.

## 10 Tye on the vagueness of consciousness

Tye (2021) argues for a “paradox” concerning the vagueness of consciousness. On the one hand, he argues that consciousness must be vague if it is identical to or metaphysically grounded in physical properties like having certain brain states or functional architectures. He does so on broadly similar grounds to those I advance above in Sect. 4 (see also Tye, 1996, 2000). On the other hand, he argues that consciousness cannot be vague, on grounds similar to those advanced by Antony and Simon: We cannot conceive of borderline cases. He attempts to resolve this paradox by arguing that consciousness has a sharp aspect, *consciousness\**. *Consciousness\** is an intrinsic property of all matter, a quiddity. Material things, in virtue of having this quiddity, are such that when they are arranged in a certain way (broadly, in the way suggested by broadcast or workspace theories) consciousness is present. But indeterminacy is still possible regarding whether an animal has or does not have the right type of cognitive arrangement, and thus ultimately in whether it is conscious.

The view I advocate here has two advantages over Tye's. First, my view – parsimoniously, attractively – does not commit to quiddities or to the panpsychist or quasi-panpsychist view that all matter has the property of being *conscious\**. I posit no new fundamental property in nature. Tye, perhaps, regards us as forced to accept *consciousness\** and its attendant quasi-panpsychism by the power of the inconceivability objection. Rather than critiquing the conceivability test, he leans upon it in developing his novel metaphysical solution. And second, in rejecting the conceivability test, my approach evades Antony's and Simon's objections. Tye's commitment to the conceivability test seems to aggravate those objections, which he does not explicitly address.<sup>23</sup>

## 11 Possible cases and applications

I conclude that the wide swath of theorists who reject panpsychism and eliminativism about consciousness and who accept a broadly naturalistic picture of the biological or functional processes that give rise to consciousness ought to accept that consciousness is a vague property admitting of indeterminate, borderline cases. It is of course possible to *tollens* that *ponens*: Theorists strongly convinced of the impossibility of indeterminate, borderline cases of consciousness (including perhaps Antony, Simon, and Goff) might prefer to reject broadly naturalistic, non-panpsychist, non-eliminativist views of the type assumed here.

---

even of all the ways in which cases might be borderline (fission, fusion, protons held somehow at intermediate distances from each other, possible quantum superpositions).

<sup>23</sup> For more discussion of Tye's view, see Schwitzgebel, 2021.

If borderline consciousness exists, where ought we expect to find it? And what implications would its existence have?

Let's start with where. Some candidates include:

- (1.) *In non-human animals of intermediate sophistication.* If we reject panpsychism, eliminativism, and saltation views, then we ought to predict that some non-human animals will be borderline conscious. Exactly which animals will depend on which theory of consciousness is true. If consciousness requires sophisticated self-representations, social mammals and birds, with good but subhuman theory-of-mind capacities, might be borderline conscious.<sup>24</sup> If consciousness requires flexible learning of the sort present in bees but (apparently) absent in simple gastropods like the sea hare, then maybe more sophisticated gastropods, like terrestrial snails and slugs, have borderline consciousness.<sup>25</sup>
- (2.) *In human fetal development.* If we assume that the fertilized egg is not yet conscious and the newborn infant is conscious, and if the empirical evidence continues to suggest no perfectly sharp, consciousness-instilling phase transition in fetal cognitive and/or neurophysiological development, then we might expect a period during which the fetus is borderline conscious. Presumably at least some neural differentiation will be required, but how far along in development such an in-between phase occurs, and how long it persists, will depend on what theory of human consciousness is correct.
- (3.) *In transit between nonconscious and conscious states.* If we assume that there are at least some moments in dreamless sleep, sedation under anesthesia, or traumatic brain injury during which people are determinately nonconscious, and if empirical evidence suggests that the transition can sometimes be gradual between these states and determinately conscious states, then we might expect borderline consciousness in some of these transit or “semi-conscious” states.

In the examples above, the entire organism is in a borderline state of consciousness. There might also be borderline conscious states in organisms that are determinately conscious at the moment in question. For example:

- (4.) *Perceptual states near the limit of conscious perceivability.* A visual stimulus presented for 200 milliseconds, then masked, will ordinarily be consciously perceived by most research participants. A visual stimulus presented for 30 milliseconds, then masked, will ordinarily not be consciously perceived, though it might influence future behavior, at least on standard views.<sup>26</sup> At intermediate presentation lengths, participants will often be uncertain whether a stimulus was presented. While it is possible that in all such cases the stimulus was either determinately experienced or determinately not experienced, another possibility is that in some such cases the stimulus is borderline conscious. For example,

<sup>24</sup> On this issue, contrast Carruthers, 2000 and Gennaro, 2012.

<sup>25</sup> Ginsburg and Jablonka, 2019, p. 395; Schwitzgebel, 2020.

<sup>26</sup> For discussion both of the standard view and challenges to it, see Peters et al., 2017. See also Brogaard's, 2010 discussion of such cases as possibly borderline.

on a theory of consciousness that ties consciousness closely to reportability, the stimulus might be borderline reportable – reportable given certain prompts or response modalities but not given others. Something similar might be possible for a tone that fades slowly into silence or a scent that slowly fades as you stroll away from its source.<sup>27</sup>

- (5.) *Perception without attention.* There’s considerable debate about the extent to which consciousness can “overflow” attention, if it can do so at all.<sup>28</sup> Do you have constant tactile experience of your feet in your shoes? Do you normally consciously experience the background hum of the refrigerator? Or do these things fade from consciousness when they fade from attention? If you’re presented with a complex visual display and given a demanding task that requires you only to attend to one part of the display, do you experience the unattended features of the display? *Maybe* every well-framed question of this sort deserves a determinate yes-or-no answer. But if borderline consciousness is possible, then some such cases might involve borderline consciousness. This might be an especially natural view if consciousness is closely linked to attention and attention is a graded phenomenon that admits of borderline cases.
- (6.) *The fringe of consciousness.*<sup>29</sup> Do you have a conscious experience of willing the actions you take? Are there a half-dozen thoughts always hanging out, so to speak, on the fringe of your awareness? Do you have constant half-conscious expectations about what is about to happen in your environment? We might consider the possibility that some such processes are borderline conscious.

I am not committing to the existence of borderline consciousness in any one type of case. What I am arguing is that the graded basis of consciousness makes it likely that borderline consciousness occurs sometimes somewhere. Absent evidence of sudden, synchronous phase changes always in exactly the right places, we ought by default expect instances of borderline consciousness wherever we can arrange a sorites series of tiny steps between cases of nonconsciousness and cases of consciousness.

If the possible existence of borderline consciousness is accepted, debates about animal consciousness, consciousness without attention, and so on, might become less stark than they currently seem. If the only options on the table are the determinate presence or determinate absence of consciousness, then it’s black versus white. With borderline consciousness, shades of gray become possible.

I leave applications to the moral status of animals and the aesthetic value of background music as an exercise for the reader.

**Acknowledgements** For helpful discussion, thanks to Michael Antony, Jonathan Birch, Liam Kofi Bright, Brit Brogaard, Philip Goff, Linus Huang, Geoff Lee, Robert Long, David Papineau, Henry Shevlin, the audience at the Kinds of Intelligence Workshop 3, and the many people who have commented on my posts on this topic on Facebook, Twitter, and my blog.

<sup>27</sup> The example of the tone fading into silence is due to Tye, 1996. Tye, 2021 rejects the example, saying that one is consciously hearing throughout, though by the end one is hearing silence – so the difference is only in the content of the auditory representation, not in the presence or absence of auditory consciousness.

<sup>28</sup> For example, Block, 2007 and the many reactions to that article; also Schwitzgebel, 2007; Prinz, 2012.

<sup>29</sup> I use “fringe” in the sense of Mangan, 2001.

**Authors' contributions** Eric Schwitzgebel is wholly responsible for all aspects of the manuscript.

**Funding** No funding was received for conducting this study.

**Data Availability** Not applicable. No data are used in this research.

## Declarations

**Competing Interests** The author has no relevant financial or non-financial interests to disclose.

**Ethics approval** Not applicable.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Alston, W. P. (1964). *Philosophy of language*. Prentice-Hall.
- Antony, M. V. (2006). Vagueness and the metaphysics of consciousness. *Philosophical Studies*, 128, 515–538.
- Antony, M. V. (2008). Are our concepts conscious state and conscious creature vague? *Erkenntnis*, 68, 239–263.
- Baars, B. J. (1988). *A cognitive theory of consciousness*. MIT Press.
- Bayne, T., Hohwy, J., Adrian, M., & Owen (2016). Are there levels of consciousness? *Trends in Cognitive Sciences*, 20, 405–413.
- Birch, J. (2021). Materialism and the moral status of animals. Manuscript dated May 24, 2021. <https://philpapers.org/archive/BIREAT-4.pdf>.
- Birch, J. (2022). The search for invertebrate consciousness. *Noûs*, 56, 133–153.
- Block, N. (1995). On a confusion about a function of consciousness. *Brain and Behavioral Sciences*, 8, 227–247.
- Block, N. (2002/2007). The harder problem of consciousness. In N. Block, *Consciousness, function, and representation*. MIT Press.
- Block, N. (2007). Overflow, access, and attention. *Behavioral and Brain Sciences*, 30, 530–548.
- Bonhomme, V., Staquet, C., Montupil, J., Defresne, A., Kirsch, M., Martial, C., et al. (2019). General anesthesia: A probe to explore consciousness. *Frontiers in Systems Neuroscience*, 13, <https://doi.org/10.3389/fnsys.2019.00036>.
- Brogaard, B. (2010). Degrees of consciousness. Manuscript at <https://sites.google.com/site/brogaardb/home> [accessed Jul. 26, 2021].
- Bruno, M. A., Vanhaudenhuyse, A., Thibaut, A., Moonen, G., & Steven Laureys. (2011). From unresponsive wakefulness to minimally conscious PLUS and functional locked-in syndromes: Recent advances in our understanding of disorders of consciousness. *Journal of Neurology*, 258, 1373–1384.
- Campbell, K. (1984). *Body and mind* (2nd ed.). University of Notre Dame Press.
- Carruthers, P. (2000). *Phenomenal consciousness*. Cambridge University Press.
- Carruthers, P. (2019). *Human and animal minds*. Oxford University Press.
- Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., et al. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. *Science Translational Medicine*, 5, 198ra105.

- Chalmers, D. J. (1996). *The conscious mind*. Oxford University Press.
- Chalmers, D. J. (2015). Panpsychism and panprotopsyism. In T. Alter, & Y. Nagasawa (Eds.), *Consciousness in the physical world*. Oxford University Press.
- Chin, C. (2015). *Borderline consciousness, phenomenal consciousness, and artificial consciousness: A unified approach* (PhD dissertation.). University of Oxford.
- Dehaene, S. (2014). *Consciousness and the brain*. Viking.
- Dennett, D. C. (1998). *Brainchildren*. MIT Press.
- Dennett, D. C. (2005). *Sweet dreams*. MIT Press.
- Frankish, K. (2016). Not disillusioned: Reply to commentators. *Journal of Consciousness Studies*, 23(11–12), 256–289.
- Gennaro, R. J. (2012). *The consciousness paradox*. MIT Press.
- Ginsburg, S., & Jablonka, E. (2019). *The evolution of the sensitive soul*. MIT Press.
- Godfrey-Smith, P. (2020). *Metazoa*. Farrar, Straus, and Giroux.
- Goff, P. (2013). Orthodox property dualism + the linguistic theory of vagueness = panpsychism. In R. Brown (Ed.), *Consciousness inside and out*. Springer.
- Goff, P. (2017). *Consciousness and fundamental reality*. Oxford University Press.
- Hall, G. (forthcoming). Is consciousness vague? *Australasian Journal of Philosophy*.
- John, E. R., & Pritchep, L. S. (2005). The anesthetic cascade: A theory of how anesthesia suppresses consciousness. *Anesthesiology*, 102, 447–471.
- Kammerer, F. (2019). The illusion of conscious experience. *Synthese*, 198, 845–866.
- Kammerer, F. (2022). Ethics without sentience. Facing up to the probable insignificance of phenomenal consciousness. *Journal of Consciousness Studies*, 29(3–4), 180–204.
- Lee, Andrew, Y. (forthcoming). Degrees of consciousness. *Noûs*.
- Lee, U. C., Müller, M., Noh, G. J., ByungMoon, Choi, & Mashour, G. A. (2011). Dissociable network properties of anesthetic state transitions. *Anesthesiology*, 114, 872–881.
- Mangan, B. (2001). Sensation's ghost: The non-sensory fringe of consciousness. *Psyche*, 7(18).
- Mashour, G. A., Roelfsema, P., Changeux, J. P., & Dehaene, S. (2020). Conscious processing and the global neuronal workspace hypothesis. *Neuron*, 105, 776–798.
- McGeer, V., & Schwitzgebel, E. (2006). Disorder in the representational warehouse. *Child Development*, 77, 1557–1562.
- McGinn, C. (1982/1996). *The character of mind* (2nd ed.). Oxford University Press.
- Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, 30, 63–134.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450.
- Oizumi, M., Albantakis, L., & Tononi, G. (2014). Giulio From the phenomenology to the mechanisms of consciousness: Integrated Information Theory 3.0. *PLOS Computational Biology*, 10 (5): e1003588.
- Overgaard, M., & Overgaard, R. (2010). Neural correlates of contents and level of consciousness. *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2010.00164>.
- Papineau, D. (1993). *Philosophical naturalism*. Blackwell.
- Papineau, D. (2002). *Thinking about consciousness*. Oxford University Press.
- Papineau, D. (2003). Could there be a science of consciousness? *Philosophical Issues*, 13, 205–220.
- Peters, M. A. K., Robert, W., Kentridge, I., Phillips, & Block, N. (2017). Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness*, 3(1), nix015.
- Prinz, J. J. (2012). *The conscious brain*. Oxford University Press.
- Roelofs, L. (2019). *Combining minds*. Oxford University Press.
- Salles, S. (2021). *Vagueness as arbitrariness*. Springer.
- Salti, M., Monto, S., Charles, L., King, J. R., Parkkonen, L., & Dehaene, S. (2015). Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. *eLife*, 4, e05652.
- Schwitzgebel, E. (1999). Gradual belief change in children. *Human Development*, 42, 283–296.
- Schwitzgebel, E. (2007). Do you have constant tactile experience of your feet in your shoes? Or is experience limited to what's in attention? *Journal of Consciousness Studies*, 14(3), 5–35.
- Schwitzgebel, E. (2011). *Perplexities of consciousness*. MIT Press.
- Schwitzgebel, E. (2014). The crazyist metaphysics of mind. *Australasian Journal of Philosophy*, 92, 665–682.
- Schwitzgebel, E. (2016). Phenomenal consciousness, defined and defended as innocently as I can manage. *Journal of Consciousness Studies*, 23(11–12), 224–235.
- Schwitzgebel, E. (2020). Is there something it's like to be a garden snail? *Philosophical Topics*. 48, 39–63.



- Schwitzgebel, E. (2021). Michael Tye on vagueness about consciousness. *Blog Post at the Splintered Mind* (Oct 13).
- Schwitzgebel, E. (forthcoming). *The weirdness of the world*. Princeton University Press.
- Seager, W. (Ed.). (2020). *The Routledge handbook of panpsychism*. Routledge.
- Searle, J. R. (1992). *The rediscovery of the mind*. MIT Press.
- Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, 8, 1391–1400.
- Simon, J. A. (2017). Vagueness and zombies: Why 'phenomenally conscious' has no borderline cases. *Philosophical Studies*, 174, 2105–2123.
- Smart, J. J. C. (1959). Sensations and brain processes. *Philosophical Review*, 68, 141–156.
- Sorensen, R. (1997/2022). *Vagueness* (Winter 2022 Ed.). Stanford Encyclopedia of Philosophy.
- Strawson, G. (1994). *Mental reality*. MIT Press.
- Strawson, G. (2006). *Consciousness and its place in nature*. Imprint Academic.
- Sy, J. L., Miao, H. Y., Marois, R., & Tong, F. (2021). Conscious perception can be both graded and discrete. *Journal of Experimental Psychology: General*. <https://doi.org/10.1037/xge0001009>.
- Tye, M. (1996). Is consciousness vague or arbitrary? *Philosophy and Phenomenological Research*, 56, 679–685.
- Tye, M. (2000). *Consciousness, color, and content*. MIT Press.
- Tye, M. (2021). *Vagueness and the evolution of consciousness*. Oxford University Press.
- Unger, P. (1988). Conscious beings in a gradual world. *Midwest Studies in Philosophy*, 12, 287–333.
- Weatherston, B. (2009). Vagueness as indeterminacy. In R. Dietz, & S. Moruzzi (Eds.), *Cuts and clouds*. Oxford University Press.
- Williamson, T. (1994). *Vagueness*. Routledge.
- Windt, J. M., Nielsen, T., & Thompson, E. (2016). Does consciousness disappear in dreamless sleep? *Trends in Cognitive Sciences*, 20, 871–882.
- Wittgenstein, L. (1951/1969). On certainty. ed. *G. E. M. Anscombe and G. H. von Wright*. Harper & Row.

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