

**When Counting Conscious Subjects, the Result Needn't Always Be a Determinate Whole  
Number**

Eric Schwitzgebel  
Department of Philosophy  
University of California, Riverside  
Riverside, CA 92521  
USA

Sophie R. Nelson  
Department of Philosophy  
New York University  
New York, NY 10003  
USA

December 12, 2024

Invited submission to a special issue of *Philosophical Psychology* on the philosophy of Daniel C. Dennett

## **When Counting Conscious Subjects, the Result Needn't Always Be a Determinate Whole Number**

Abstract: Could there be  $7/8$  of a conscious subject, or 1.34 conscious subjects, or an entity indeterminate between being one conscious subject and seventeen? Such possibilities might seem absurd or inconceivable, but our ordinary assumptions on this matter might be radically mistaken. Taking inspiration from Dennett, we argue that, on a wide range of naturalistic views of consciousness, the processes underlying consciousness are sufficiently complex to render it implausible that conscious subjects must always arise in determinate whole numbers. Whole-number-countability might be an accident of typical vertebrate biology. We explore several versions of the inconceivability objection, suggesting that the fact that we cannot imagine what it's like to be  $7/8$  or 1.34 or an indeterminate number of conscious subjects is no evidence against the possibility of such subjects. Either the imaginative demand is implicitly self-contradictory (imagine the one, determinate thing it's like to be an entity there isn't one, determinate thing it's like to be) or imaginability in the relevant sense isn't an appropriate test of possibility (in the same way that the unimaginability, for humans, of bat echolocation experiences does not establish that bat echolocation experiences are impossible).

Word Count: ~8000 words, plus 4 figures

Keywords: Artificial Intelligence; consciousness; Dennett, Daniel; personal identity; split-brain; subjectivity

# When Counting Conscious Subjects, the Result Needn't Always Be a Determinate Whole Number

## 1. Introduction

People typically assume that conscious subjects come in discrete, countable wholes. Either there are no conscious subjects in the seminar room, or there are two, or six, or twenty-three. It seems bizarre to say that there might be 9.382 subjects, or  $7/8$  of a subject, or  $3i$  subjects, or that the number of subjects might be indeterminate between five and seventeen or best represented by a nine-dimensional non-Euclidian surface. What could such remarks even mean, unless as a wry comment about a shy student with 30% of their head poking through the door?

We will argue that this ordinary assumption is false. When counting conscious subjects, the result needn't always be a single or determinate whole number. Ordinary counting might fail, or yield a mathematical representation other than a whole number, or yield a set of whole numbers among which the result is indeterminate. At best, discrete countability is an accident of typical vertebrate biology, which needn't apply to invertebrates, AI systems, or atypical humans. If it's part of our concept of consciousness that conscious subjects necessarily come in determinate, whole bundles, then our ordinary concept of consciousness requires repair.<sup>1</sup>

We start by clarifying our notion of a conscious subject. After arguing that the complexity of the processes underlying consciousness renders our thesis plausible, we argue for the possibility of cases indeterminate or intermediate between zero subjects and one, as well as the possibility of cases indeterminate or intermediate between one subject and whole numbers greater than one. We conclude by addressing objections concerning inconceivability.

---

<sup>1</sup> For similar views, see Nagel 1971; Bostrom 2006; Fekete, Van Leeuwen, and Edelman 2016; Roelofs 2019, in draft; Lycan 2022; Schwitzgebel & Nelson 2023; Salisbury 2023. For objections, see Schechter 2018 (pp. 19-24); Brook & Raymond 2001/2017.

Our position is partly inspired by Dennett, and we believe he would have endorsed it. According to his “Multiple Drafts” theory of consciousness, not only do different “drafts” of experience exist at different times (generating his famous “Orwellian” vs. “Stalinesque” dilemma); they also exist simultaneously. He writes, “At any point in time there are multiple drafts of narrative fragments at various stages of editing in various places in the brain” (1991, p. 135). The narratives that constitute our “selves” “issue forth *as if* from a single source” encouraging us “to (try to) posit a unified agent” (1991, p. 418, emphasis in original). While discussing split-brain and multiple personality cases, Dennett says that we typically want “to preserve the myth of selves as brain-pearls, particular, concrete, countable things, rather than abstractions... refusing to countenance the possibility of quasi-selves, semi-selves, transitional selves” (1991, pp. 424-425).<sup>2</sup>

## 2. *Our Main Claim*

Marcel sips tea containing the crumbs of a madeleine cookie and reflects back on his days in Combray. In that moment, he simultaneously experiences the warmth of tea in his mouth, the taste of the madeleine, and recollections of Combray. These three experiences (or experience-parts or experience aspects) all belong to the same field of consciousness. They are *phenomenally unified*. There’s something it’s like to have them together. Perhaps, as Bayne and Chalmers (2003) suggest, they are subsumed within a more complex experience of tea-with-madeleine-with-recollection. Perhaps, as Dainton (2000) suggests, they’re united by a basic, not-

---

<sup>2</sup> Despite such remarks, one might interpret his “illusionism” about consciousness as implying that there is always determinately zero conscious subjects. See Dennett 1991 and 2016. For a discussion of tensions between Dennett’s seemingly anti-realist and seemingly realist claims about consciousness, see Schwitzgebel 2007 and Dennett’s 2007 reply.

further-analyzable relation of co-consciousness. We take no stand on how to analyze phenomenal unity, but we assume that the core idea is clear: two experiences might belong either to the same, or to different, fields or streams of conscious experience.<sup>3</sup>

What does it mean to be determinately one conscious subject? *Conscious subjects*, as we'll understand them here, are just such bundles of experiences, individuated by maximal relations of phenomenal unity. If experience A is unified with experience B, and if experience B is unified with experience C, and if experience A is unified with experience C, then there is one conscious subject with a unified experience of A-with-B-with-C.<sup>4</sup> If experience D is unified with no other experiences, then there is one conscious subject experiencing only D.

This notion of a “conscious subject” is unusual but not unprecedented.<sup>5</sup> We choose it not because we endorse it as the best approach to subjecthood from the point of view of theories of selfhood or personhood, or because we think that “subjects” on this conception are identical to cognitive systems, bodies, or persons, but rather because it targets a specific phenomenon of special interest. Broader definitions of “subject” that appeal to more than conscious experiences at a moment—for example, diachronic relations between conscious states; or causal threads of

---

<sup>3</sup> For overviews, see Brook & Raymond 2001/2017; Dainton 2014; Hill 2018; Masrour 2020.

<sup>4</sup> We speak here as though experiences within unified fields of consciousness are themselves countable. However, when we observe our own phenomenology, it seems as though there's no one right way to divide it into pieces. See also Bayne 2010 and Builes 2021. For an early attempt to discover the minimal elements of conscious experience, see the “elementism” of Wundt 1897/1897. Searle 2000 and Tye 2003 argue for a trivial version of the view that unified fields of consciousness contain a determinate number of experiences: always exactly one.

<sup>5</sup> For instance, Strawson 2003 describes “thin” subjects that aren't ontologically distinct from the unified set of experiences they possess (similarly Yetter-Chappell in draft). Bayne 2010 (Ch. 12) characterizes subjects as “intentional entities” each necessarily associated with one unified stream of consciousness. Shani & Williams 2022, drawing on Coleman 2013, analyze subjects merely as unified perspectives containing unique, partially ordered sets of phenomenal properties.

memory and inference; or values, personality traits, and nonconscious mental processes; or bodies—probably make it easier to defend our thesis. Temporally extended views, for example, appear to permit Parfitian (1984) cases of merging, splitting, and slow change. Views that rely on synchronous nonconscious processes appear to permit hypothetical AI or alien cases of synchronous partial unity or overlap due to shared nonconscious mechanisms.<sup>6</sup> Similarly, bodies can plausibly overlap. By focusing on “conscious subjects” in our narrow sense, we address the countability issue in its starkest form, where skeptical interlocutors’ intuitions seem to run strongest against our view.

On this notion of a “conscious subject”, is at least conceptually possible for multiple conscious subjects to exist in the same body (one experiencing A-with-B-with-C, another experiencing D-with-E-with-F), perhaps even coordinating to operate as a single person.<sup>7</sup> On some accounts of personal identity, *you* might be nothing but a conscious subject, perhaps extended over time. But even if you aren’t strictly identical to a conscious subject, you are intimately related to one. If you are, right now, experiencing the warmth of tea, the taste of madeleine, and recollections of Combray, then there is a conscious subject in your spatiotemporal vicinity undergoing exactly that set of token conscious experiences, and that conscious subject is certainly not someone other than you. Perhaps it is a part of you, an aspect of you, a process within you, or in some other way ontologically adjacent.

Our thesis is that conscious subjects, in the sense just articulated, needn’t always come in determinately whole-number units. Experiences don’t always bundle neatly.

---

<sup>6</sup> See Schwitzgebel & Nelson 2023

<sup>7</sup> For example, see Schechter 2018 for an argument that split-brain patients are two conscious minds in one body.

### 3. *Plausibility Considerations*

For the purposes of this argument, we assume a broadly naturalistic approach to consciousness on which human beings have conscious experiences in virtue of something about the complex structure and composition of their bodies and brains. (We don't assume materialism. Some forms of neutral monism, property dualism, and even panpsychism are naturalistic in the relevant sense.) The correct naturalistic theory of consciousness is far from settled. However, on a wide range of plausible theories, it appears possible to construct cases that defy whole-number countability, even if such cases are rare among vertebrates.

Consider Global Workspace theories, for example, according to which information is conscious if it is broadcast through a "global workspace" for potential use by downstream cognitive systems.<sup>8</sup> How large must a workspace be, and how much information must it be capable of feeding downstream, before it is appropriately "global"? Is there some minimum workspace size or quantity of connectedness below which an entity is not at all conscious and above which an entity is conscious, with no intermediacy or indeterminacy? Could workspaces be partial or overlapping? If information is available for 85% of downstream processes, is that enough? 51.578103%? Two? Imagine an architecture with downstream processes in two clusters, A1, A2, and A3 vs. B1, B2, B3, and B4. What if some information is available to all of A and B while some is available only to A or only to B or only to A1 and B3? Must there be a bright line between cases of a single unified conscious subject with imperfect downstream information uptake and two discrete subjects with limited or indirect information sharing? The model seems to allow in principle for a diverse range of messy cases where subjecthood and unity would

---

<sup>8</sup> Influential formulations include Baars 1988 and Dehaene 2014.

either be partial or best captured with more complex mathematical representations such as vectors or regions.

Consider Higher-Order theories of consciousness, according to which an experience is conscious if it's the target of a higher order representation of the right sort.<sup>9</sup> Perhaps in typical vertebrate cases, there's a single higher-order system that does all of the tracking and unifies all of the represented lower-order states. But as with Global Workspace theory, we can imagine considerably messier architectures. The same lower-order state might be targeted by different higher-order systems that link to different or partially overlapping output systems. Partly overlapping sets of lower-order states might be targeted by different higher-order systems that are partly integrated and have different downstream influences on the organism. It might be only in cases at the extremes that there exists either one unified conscious subject or many sharply distinct subjects. The majority of possibilities might lie in the middle.

Consider Dennett's (2005) "fame in the brain" view, according to which a cognitive process is conscious to the extent that it is "famous"—that is, influential on—other cognitive processes. Fame among humans is a complex phenomenon. People can be famous in some circles, unknown in others. Fame circles can nest and overlap. Fame relationships can correlate imperfectly, so that if Person A is known to Person X, they are 85% likely to be known to Person Y but only 10% likely to be known to the average member of the population. While it seems to be typical in the human case that if a process—say a representation of the color red in some region of the visual field—is "famous" in the centers that govern verbal report, it will also be famous in the centers that govern long-term planning and control of the fingers. But widespread

---

<sup>9</sup> For a review, see Carruthers & Gennaro 2001/2020. Rosenthal 2003 provides an account of unity of consciousness on a higher-order theory.



dissociations do occur—for instance, in split-brain patients.<sup>10</sup> As Dennett (1991) emphasizes, dissociations even occur in ordinary humans. This suggests a complex range of partial cases that defy characterization as exactly “one conscious subject” or two, or three, or seventeen.

Similar considerations apply to virtually all naturalistic theories of consciousness. Almost inevitably, they ground consciousness in processes or structures that can be implemented in fuzzy, indeterminate, or partly overlapping ways, suggesting mechanisms for unity that can be implemented in fuzzy, indeterminate, or partly overlapping ways. Embodied theories of consciousness are no exception, despite the fact that we usually imagine embodied organisms as discrete individuals. Sponges, lichen, grasses, and birch forests connected at the root are not as readily divisible into discrete individuals, determinately countable. Octopus arms operate partly independent of the head: is there one conscious subject or nine?<sup>11</sup> The Hogan Twins have overlapping brains and differ in personality while having the capacity to report each other’s sensory experiences.<sup>12</sup> AI architectures might similarly admit of various types and degrees of integration. Must it really always be the case, for every possible pair of craniopagus twins, that there is either determinately one conscious subject or two?

In general, if consciousness has a complex, naturalistic basis, it seems correspondingly plausible that cases could arise where the number of conscious subjects is similarly complex, defying determinate whole-number countability.<sup>13</sup> Those who would argue otherwise owe us an

---

<sup>10</sup> For reviews, see Volz & Gazzaniga 2017 and Schechter 2018.

<sup>11</sup> See for example Godfrey-Smith 2016 and Carls-Diamante 2017.

<sup>12</sup> See the 2017 CBC documentary *Inseparable: Ten Years Joined at the Head*; Dominus 2011; Cochrane 2020; Kang 2022.

<sup>13</sup> One exception is Integrated Information Theory (Oizumi, Albantakis & Tononi 2014), the mathematics of which will always generate a whole number of conscious subjects. As Schwitzgebel has emphasized in other work (Schwitzgebel 2014, 2024), this aspect of IIT creates implausibly sharp lines, such that tiny structural differences can constitute arbitrarily large differences in the number of conscious subjects.

explanation of why conscious subjectivity must always be determinately countable despite the complexity and multidimensionality of the processes that generate it.

#### **4. *Between Zero and One***

Start with a system that is not and does not have a brain. Posit a series of changes that render it, eventually, an apparently conscious system that determinately is or that determinately has a brain. Maybe the system begins as a fertilized embryo and ends as a newborn child. Maybe it begins as a simple form of artificial intelligence and ends as an extraordinarily sophisticated one. Evaluate the changes at arbitrarily precise time scales. Theories of conscious subjecthood then face a quadrilemma. Either (1) the system is a conscious subject the whole way through, even when it lacks a brain entirely; (2) the system is never a conscious subject; (3) there is a sharp, stepwise distinction where the system suddenly becomes a conscious subject; or (4) there is a non-sharp or non-stepwise distinction between when the system is not a conscious subject and when it is.

For any particular transformation of this sort, option (1) will always be available to certain kinds of radical panpsychists, who hold that conscious experience is present not only in all fundamental entities but also in all composites of fundamental entities. Such a panpsychist could say that the collection of fundamental particles that constitutes the initial, non-brain system somehow already possesses a unified field of conscious experience. Similarly, option (2) will be available to radical eliminativists, who hold that conscious subjects don't exist at all.<sup>14</sup> However,

---

<sup>14</sup> Note that not all panpsychist positions hold that literally *everything* is conscious, and consequently, not all panpsychists can take option (1). For example, as Goff (2013) notes, panpsychists who hold that fundamental particles and humans are conscious, but not rocks and most other composite entities, still face the dilemma between (3) and (4). Note also that not all "eliminativist" or "illusionist" positions hold that nothing is conscious. Frankish 2016, for

for present purposes, we assume that options (1) and (2) won't apply universally to all such gradual series. We must then sometimes choose between (3)—saltation—and (4)—indeterminacy or intermediacy. Note that intermediacy differs from indeterminacy. The number of conscious subjects could be intermediate (for example, exactly 0.3589) without being indeterminate. Conversely, the number could be indeterminate (not determinately zero, not determinately one) without being intermediate.

Saltation is unattractive on broadly the grounds discussed in Section 3. It would be surprising if, in general, atop smooth gradations of gradual structural change, conscious subjecthood suddenly jumps in, with no indeterminacy or intermediacy at all. Are we to imagine, for example, that lizards of one genus determinately are conscious subjects while lizards of another genus, the tiniest bit less sophisticated, are not? Nature does admit of sudden phase transitions. Water freezes at exactly  $0.0^\circ$ . Beams suddenly snap under loads. But, except in quantum cases, even such phase transitions aren't perfectly sharp. Close inspection reveals intermediate states. Must there always be an exact moment, down to the millisecond, down to the nanosecond, at which conscious subjecthood suddenly pops in, with no intermediate or indeterminate phase? Furthermore, must *all* such sequences lack *any* intermediate or indeterminate phases? This is a bold claim! Only a powerful argument could justify it, and we are aware of no published arguments within the framework of scientific naturalism that attempt to do so.<sup>15</sup>

---

example, holds that nothing is “phenomenally conscious” in a certain technical sense of the phrase while allowing that some entities are conscious in a weaker sense of “consciousness” is understood in our intended sense, stripped of dubious theoretical commitments.

<sup>15</sup> As discussed in Schwitzgebel 2023, there are some arguments specifically against indeterminacy concerning the presence or absence of consciousness (e.g., Antony 2008; Goff 2013; Simon 2017), but it's unclear whether these arguments generalize to determinate intermediacy.

Accepting that the number of conscious subjects can be indeterminate or intermediate between zero and one doesn't settle the best numerical means for representing that possibility. The best approach in some cases might be a real value between zero and one. For example, if we knew that a global workspace of size 1000 constituted exactly zero conscious subjects, a global workspace of size 2000 constituted exactly one conscious subject, and all intermediate workspace sizes were determinately intermediate between zero conscious subjects and one, we could use values between 0 and 1 to represent the number of subjects present. (It might remain open whether the number of subjects related linearly, logarithmically, sigmoidally, or in some other way to the size of the workspace.) Similarly, a real value could measure how close a case of indeterminacy lies to the border of determinacy (e.g., an indeterminate case much closer to being determinately one conscious subject than zero might be represented by the number 0.9). On the other hand, real numbers might suggest implausible precision, unless construed as mere approximations. A less commissive representation might simply be the open interval  $(0,1)$  or a set of possibilities among which the number is indeterminate  $\{0, 1\}$ . Or one might opt for more structure: a vector or region that represents several independent dimensions of intermediacy or indeterminacy. Even imaginary numbers might be applicable if consciousness involves quantum states represented by complex numbers.

##### **5. *Between One and N***

Luke Roelofs (2019) constructs a similar slippery-slope case between two subjects and one. Start with two conscious brains, wholly distinct. Slowly join them, one neural connection at

a time, until they form a single, unified subject of experience.<sup>16</sup> Either (1) the same, non-zero number of conscious subjects is present the whole way through; (2) there aren't, in fact, any conscious subjects during any stage of the process; (3) there's a sudden, sharp change in the number of conscious subjects; or (4) there's a gradual, non-sharp change in the number of conscious subjects.

Again, we assume the falsity of (1) and (2), although (1) is available to radical panpsychists while (2) is available to radical eliminativists.<sup>17</sup> Option (3) seems just as implausible here as in the zero-to-one case. If each step of integration is sufficiently tiny, the architectural and functional differences will be correspondingly tiny, and it's unattractive to suppose that the seemingly huge metaphysical difference between one and two conscious subjects would be grounded in a tiny architectural or functional difference. And recall that the saltationist is committed to a negative universal: There cannot be *any* way of slowly connecting *any* two conscious subjects such that there is a *single moment* of indeterminacy or intermediacy. Such a bold claim requires compelling support.

If  $N > 1$ , the best numerical representation of indeterminacy or intermediacy is unlikely to be a single real number value. For example, in earlier work (Schwitzgebel & Nelson 2023), we describe a case that allows a slippery slope between 1 and 201. Imagine a conscious AI, perhaps employing a futuristic technology very different from silicon chips. The entity is composed of a large, orbiting AI system plus 200 robotic bodies on a planetary surface, each

---

<sup>16</sup> This thought experiment is closely related to gradual brain bisection scenarios used for similar gradualist ends in Nagel 1971 and Lockwood 1994. Further precedents appear in Hirstein 2012 and Sotala & Valpola 2012.

<sup>17</sup> We should note that Roelofs accepts horn (1) of both quadrilemmas. Although Roelofs holds that there is no discrete moment at which two “intelligent subjects” become one, Roelofs embraces radical panpsychism regarding *conscious subjects* in our sense, holding that every mereological sum of concrete entities constitutes a distinct conscious subject.

with their own local AI processors. If the entity is massively interconnected in the right way, it is plausibly a single conscious subject with multiple bodies or one spatially discontinuous body.<sup>18</sup> If the entity is sparsely connected, or connected in the wrong way, there are plausibly 200 or 201 distinct conscious subjects in communication with one another (201 if the orbiting system is conscious, 200 if not). Between these extremes lies approximately a continuum of possibilities, some of which may constitute cases indeterminate or intermediate between 1 and 201. However, it would be misleading to numerically represent the number of conscious subjects in an AI system indeterminate between 1 and 201 in exactly the same way we represent the number when 101 typical humans gather in a room. Furthermore, such cases can be structured to involve multiple independent dimensions of intermediacy or indeterminacy, depending on the degree of integration between the orbiting AI system and each individual robot, thus inviting multidimensional mathematical representation.

We can also construct a *zero to N* case as follows: Start with our orbiter and robots, all simple enough that none is conscious. Improve them all simultaneously, one arbitrarily small step at a time. By analogy with our non-saltationist reasoning above, there should be a range of cases indeterminate or intermediate between zero and N.

## 6. “*But It’s Inconceivable!*”

As we descend further into such unfamiliar ways of thinking, we expect that many readers will stop somewhere short, with a worry along these lines: It is inconceivable, and therefore metaphysically impossible, that conscious subjects exist in anything other than

---

<sup>18</sup> For other examples of spatially distributed conscious subjects, see Dainton 2000; Bayne 2010; Schwitzgebel 2015, 2024.

determinately countable wholes. At least, some readers may be concerned that conscious subjects lacking determinate or discrete countability are not “positively” conceivable, constituting grounds to doubt that indeterminate or intermediate cases of conscious subjects are metaphysically possible.<sup>19</sup> In what follows, we’ll consider three distinct inconceivability challenges for our view.

First, however, an observation. In conversation, we find philosophical interlocutors typically more open to indeterminacy than to intermediacy or other types of non-whole-number representation, and to express particular skepticism about fractional subjects, such as  $7/8$  of a subject. Perhaps they are right to be skeptical. It’s not obvious what could make a cognitive system constitute some *particular* fractional number of subjects (e.g., 0.8) rather than some other fractional number. Here, we remind the reader that we argue only for the disjunctive thesis of indeterminacy *or* non-whole-numbers. That said, we have never seen a well-developed *argument* against fractional subjects from a (non-panpsychist, non-eliminativist) naturalist perspective.<sup>20</sup>

On non-whole-number representations larger than one, imagine a cloudscape. Sometimes clouds are straightforwardly countable. Often, however, complex cloudscape are best represented in mathematically complex ways. Geometrical representations might enable accurate representation of how the clouds clump and cluster, perhaps with high mathematical precision. One needn’t shruggingly say only “it’s indeterminate but more than three”. Similarly, one non-


---

<sup>19</sup> Our use of “positive conceivability” follows Chalmers 2002. On relations between conceivability and possibility, see also Gendler & Hawthorne 2002; Mallozzi, Vaidya & Wallner 2007/2023.

<sup>20</sup> Views that treat phenomenal unity as indeterminate—rather than determinately intermediate—face one objection that views allowing only intermediacy do not. The following propositions form an inconsistent triad: (1) There’s no vagueness at the fundamental level of reality; (2) Phenomenal unity is fundamental; (3) Phenomenal unity can be vague. Many have argued for (1) (e.g., Heller 1996); some have argued that (2) true is true as well (e.g., Roelofs 2019). If these arguments are successful, then (3) must be false. Dennett would have denied (2).

whole number possibility worth considering is a massive AI system with unity and disunity relations best understood by complex network representations.

### 6.1. Between Zero and One



**Figure 1:** An entity intermediate or indeterminate between zero and one conscious subjects. Dotted circles represent indeterminately or partially conscious mental states.

To conceive of a case intermediate or indeterminate between zero and one, we must conceive of indeterminately or intermediately conscious mental states (see Figure 1). After all, if there's a single, wholly conscious mental state present, there's at least one conscious subject. Can we imagine indeterminate or intermediate consciousness? Return to Marcel's experience of warm tea, a taste of madeleine, and recollections of Combray.

Now, subtract the experience of warm tea and the taste of madeleine. Marcel's mental state is not one-third conscious, and consequently, Marcel has not been reduced to one third of a conscious subject. He's still a conscious subject, just with fewer experiences. Now, subtract Marcel's recollections of Combray. No conscious state remains, and therefore, no conscious subject does either (at least at that moment, according to our definition). This constitutes a discrete jump from one to zero. Could we somehow subtract approximately *half* the experience of recollection, making Marcel somehow *half* conscious? We might imagine first a vivid recollection of his bedroom in Combray with the glimmering flame of the nightlight in its bowl of Bohemian glass, hung by chains from the ceiling; and then we might reduce the experience. Forget the chains, forget the glass, recall only the glimmering flame. Still, that memory seems determinately to be



an experience, and thus, one conscious subject determinately exists. Perhaps the flame can be remembered only vaguely—what was its shape? Its color? Was it even a flame, or only a glowing coal? However faint we imagine Marcel’s phenomenology, it seems we still have some phenomenology, until we remove the experience altogether and thus conscious subjecthood altogether. No intermediate state seems conceivable. Either there’s something it’s like to be Marcel in that moment—no matter how simple—or there’s nothing it’s like. No conceptual space stands open between something and nothing. A half-something is already a something, unless it is nothing.

We acknowledge the pull of this way of thinking. However, we reply that there’s an implicit self-contradiction in any attempt to imagine what it’s like to be an entity indeterminate or intermediate between zero and one. Necessarily, there’s no one determinate thing it’s like to be a borderline case of a conscious subject. No one, regardless of their cognitive architecture, can grasp what it’s like to be such an entity any more than they can grasp what it would be like to perceive a square circle. The more you try to vividly imagine the phenomenology of something that, by nature, lacks even a single wholly conscious, determinate experience, the worse you miss the mark.

That said, entities who regularly enter intermediate or borderline conscious states might have no trouble discerning or self-representing such states, and they might develop corresponding concepts. For example, they might be able to think, “I spend a considerable period of time in that *mizzy* state between being neither determinately unconscious nor determinately conscious” (alternatively, partly conscious, if intermediacy is possible). They might be even able to imaginatively reconstruct *mizzy* states through imaginative episodes that are themselves *mizzy*. Nonetheless, just like us, they would fail to conceptualize the full, determinate experience of


borderline or partial conscious subjecthood, because there *is* no full, determinate experience of borderline or partial conscious subjecthood.<sup>21</sup>

## 6.2. *Between One and N*

What would be required to conceive of an indeterminate or non-whole-number of subjects between one and  $N$  (where  $N$  is a whole number greater than one)? We see two possibilities between which we remain neutral.

### 6.2.1. *Indeterminate or intermediate unity among determinately conscious experiences.*

First, in order to imagine subjects indeterminate between one and  $N$ , one might attempt to



**Figure 2:** An entity intermediate or indeterminate between one and three conscious subjects. Solid circles represent determinately conscious mental states. Dotted lines represent indeterminate or intermediate unity among those states.

imagine *determinately* conscious experiences that are *indeterminately* or *intermediately* phenomenally unified. Rather than determinately one or several distinct, unified bundles of conscious experience—for example, A-with-B-with-C and D-with-E-with-F—we might conceive of something in-between (see Figure 2). Perhaps the relation of phenomenal unity admits of degrees, or there can be cases in which its presence is vague. On this picture, if

Marcel were a cognitive system intermediate or indeterminate between one and two subjects, he would need to have both a determinate experience of the warmth of tea and a determinate experience of the taste of madeleine, without those experiences being either fully (that is,

---

<sup>21</sup> For a more detailed treatment of the Paradoxical Demand reply to the case of indeterminate consciousness, see Schwitzgebel 2023.


determinately and not intermediately) unified into a conscious whole nor fully disunified into two separate wholes. He would need to feel both experiences simultaneously, but *neither entirely together nor entirely apart*.

Trying to conceive of what that would be like, we might imagine the experience of the warmth phenomenally unified with a hazy, faint, or peripheral-seeming experience of the taste (or vice versa). But this isn't a case of warmth indeterminately or intermediately unified with the more robust taste from the original scenario; it's just a case of warmth determinately phenomenally unified with a hazy, faint, or peripheral-seeming taste—a determinately conscious but less vivid experience. A worry potentially arises: However much we try to imagine conscious experiences being only partially severed from one another, it seems we're always left imagining wholly unified experiences, and therefore a single conscious subject—that is, until we sever the tie altogether and are left with determinately multiple conscious subjects. We imagine Marcel having his experiences conjointly, or we imagine the experiences being felt in isolation. Those are, it seems, the only conceivable options.

But if our hypothetical objector seeks an act of imagination that joins together, into the objector's own, fully unified imaginative experience, two experiences that aren't fully unified, then again make an implicitly self-contradictory demand. Of course, there isn't one fully unified thing it's like to feel two indeterminately or intermediately unified experiences simultaneously. If the experiences were determinately *disjoint*, the incoherence of this imaginative demand would be obvious: We can't expect to imagine what it's like to have Marcel's experience of the warmth of tea jointly with Odette's experience of a piano sonata. Clearly, there's no one thing it's like to have these two disjoint experiences. The same holds, though less obviously, when the two experiences are indeterminately or intermediately conjoint.

This isn't to say that indeterminate or intermediate phenomenal unity is inherently unimaginable. For example, an indeterminately or intermediately phenomenally unified conscious subject *with an indeterminately or intermediately unified imagination* might have no trouble introspecting and reflecting on such mental states. They might be able to think, "When my tentacles with neurons at their tips are only partly connected to my body, I start having *dissy* experiences that aren't fully unified with one another." Although they couldn't imagine *dissy* experiences as being like *one* thing any better than we can, they might be able to conceive of *dissy* states through imaginative episodes that are themselves *dissy*. Suppose Marcel experiences the warmth of tea in one tentacle and experiences memories of Combray in another tentacle, and these experiences are only partly unified. Odette—or Marcel himself later—might imagine or remember this phenomenology by means of an imaginative structure that is itself distributed among different tentacles and not fully unified. This would not (we assume) be much like normal human imagination, but what is alien to us needn't be impossible.

6.2.2. *Intransitive unity.* Alternatively, subjects might be intermediate or indeterminate between one and N because they have overlapping fields of consciousness that share token experiences. This would require the relation of phenomenal unity to be *intransitive*. Most philosophers have assumed that if an experience A is unified with another experience B, and if B



**Figure 3:** An entity intermediate or indeterminate between one and two conscious subjects, due to intransitive unity relations, despite determinately conscious mental states and determinate relations of full unity among subsets of states.

is unified with a third experience C, then A and C are themselves unified. However, if intransitive unity—or, as many have called it, partial unity—is possible, then A and B might be unified, and B and C might be unified,

without A and C being unified as well (see Figure 3).<sup>22</sup>

To understand how subjects intermediate or indeterminate between one and N would look if phenomenal unity were intransitive, first consider two cases: Tiny Overlap and Massive Overlap. In Tiny Overlap, two alien brains (or AI systems) share a tiny bit of tissue. At the target moment, each alien has a million experiences distinctive to them—vast turbulences of vivid, unified experiential activity—plus one tiny shared experience: the faint sound of a distant motor. Plausibly, Tiny Overlap should be conceptualized as a case of two conscious subjects who happen to share one token experience. Massive Overlap is the complementary case: a million experiences are shared, but two experiences are not—a green dot in the left visual periphery and a red dot in the right visual periphery. The green dot is unified with everything but the red dot, and the red dot is unified with everything but the green dot. Plausibly, Massive Overlap should be conceptualized as a case of a single conscious subject who is ever-so-slightly disunified. Now, imagine that, like in Section 5, Tiny Overlap gradually becomes Massive Overlap as the aliens’ brains slowly fuse. In the middle of this process, there would be two partially overlapping fields of experience that would plausibly count as neither determinately one nor two subjects.

Again, intransitive unity might seem impossible because we’re unable to imagine what having an intransitively unified perspective would be like.<sup>23</sup> If we try to imagine how it would feel to experience the warmth of tea alongside the taste of madeleine, and the taste of madeleine

---

<sup>22</sup> For defenses of intransitive phenomenal unity, see Lockwood 1989; Tye 2003; Schechter 2014; Salisbury 2023; Yetter-Chappell in draft. For objections, see Dainton 2000; Hurley 2003; Bayne 2010; Vogel 2014. For discussion of the closely related view that numerically distinct subjects can share numerically identical experiences, see Hirstein 2012; Roelofs 2016, 2019; Cochrane 2020; Goff & Roelofs forthcoming. Although the phrases “partial unity” and “intransitive unity” have previously been treated as synonymous, “partial unity” might better be reserved for the genus of which the cases described in 6.2.1 and 6.2.2 are species.

<sup>23</sup> See Dainton (2000) and Bayne (2010) for versions of this objection.

alongside recollections of Combray, while the warmth of the tea and the recollections of Combray remain disunified, we draw a blank. Once again, however, the demand is paradoxical. This hypothetical objector seeks to imagine, in a transitively unified way, what it would be like to have intransitively unified experiences. They seek to bring into their unified imagination an experience of A with B with C that is somehow simultaneously an experience of A with B and B with C but not A with C. This self-contradictory imaginative standard is inappropriate to the case.<sup>24</sup> Intransitively unified experiences aren't like any *one* thing, or part of any *one* phenomenal perspective.

Still, we see no reason to believe that a creature with intransitively unified consciousness couldn't imagine, *in an intransitively unified way*, intransitively unified experiences. Maybe Marcel's Arm 1 is unified with his head, and his head is unified with Arm 6, but Arms 1 and 6 are not unified. Even though Marcel can't join the experiences of his head and arms in one phenomenally unified act of imagination, he might imagine, think about, introspect, or recall them by means of a similarly distributed imagination.

Our response to each of the three conceivability objections is the same. The objections, we suspect, turn on our wanting to imagine these extraordinary cases in the ordinary, unified way. It is of course often fine to want to imagine things in the ordinary way. That's why the imaginative demands superficially seem reasonable. Sometimes we feel like we can't quite get our head around an experience if we can't imaginatively construct or recreate it in the way we imaginatively construct or recreate an experience of seeing a sunlit mound of gold while feeling its warm, smooth texture. The ordinary way we imagine experiences is as determinate

---


<sup>24</sup> For a similar reply to the inconceivability challenge facing the view that phenomenal unity can be intransitive, see Dainton 2000 (Ch. 4, p. 98) and Schechter 2014.

experiences (perhaps with some indeterminate *contents*) that are wholly, determinately, transitively unified. It is inappropriate and paradoxical to apply this standard of imagination to the non-whole-number cases at hand.

More reasonably, one might alter the standards of imagination to appropriately match the target cases. *Mizzy*, *dissy*, and intransitive states might seem alien to us, but they are unimaginable only in the same way that color is arguably unimaginable to people blind from birth and bat echolocation seems unimaginable to ordinary humans. Human limitations, not metaphysical or nomological impossibility, explain our inability to concretely imagine conscious subjects that don't come in determinate, countable whole numbers.

### 6.2.3 Countability partly recovered?

If we accept the intransitivity approach to conceiving of subjects indeterminate or intermediate between one and N, there still might be a sense in which fields of experience are countable. After all, overlapping entities can come in determinate whole numbers.  $A = \{x \mid 1 \leq x \leq 999,999\}$  and  $B = \{x \mid 2 \leq x \leq 1,000,000\}$  aren't indeterminate between two sets; rather,



**Figure 4:** Two regions of space, determinately countable despite massive overlap.

they're determinately two sets that share 99.9999% of their contents. The shapes in Figure 4 don't occupy somewhere between one and two regions of space. They simply occupy massively overlapping regions. If it's possible in principle to measure or quantify the degree of overlap between two fields of consciousness, it might be possible in principle to count up a whole number of

overlapping conscious fields, just as it's possible to count up a whole number of overlapping sets and an overlapping number of bounded, continuous regions of space.

One might worry that understanding subjects this way would once again force us to posit a sharp, inelegant jump atop a complex physical base. However, we see no reason for concern: This approach would still allow for gradual phenomenological transitions alongside gradual increases and decreases in functional integration. Consider again an analogy to sets: suppose I pull out a piece of paper and list the sets of numbers [1, 2, 3] and [2, 3, 4]. I now have two mathematical objects instantiated on my paper. Suppose that I then erase the number one from the first set and the number four from the second set. Now, I only have one mathematical object instantiated. We could, at this point, be puzzled – how bizarre that the transition from one entity to two was so clean! But abstract objects can be distinct, even when the differences between them are tiny, and they can collapse into one when those tiny differences are erased. The same goes, potentially, with “conscious subjects” if we define them in terms of overlapping sets or fields of experience.

However, generalizing this approach to counting subjects generates what would seem to be intuitively, pragmatically, and functionally the wrong result in the Massive Overlap case: The tiniest bit of intransitivity would generate two discretely different “subjects” even if the system is functionally, practically, and introspectively almost identical to a single subject. This would be especially worrying if, as Dennett suggests, ordinary humans are themselves often not fully unified. Even if it's possible in principle to count up massively overlapping fields of experience, calling each one a “subject” seems bizarre. Furthermore, this restoration of countability, if it works at all, works only for the intransitive unity case. Zero-to-one and indeterminate or intermediate unity cases will still ruin any strict countability principle.



In Section 2, we defined “conscious subjects” as bundles of experience, individuated by maximal relations of phenomenal unity. If, as we’ve just suggested, unity can come in degrees or be intransitive or indeterminate, the relevant relation might be *sufficient* rather than *all-or-none* unity and there might not always be a single best individuation scheme.

## 7. Conclusion

Dennett challenged readers to leave behind their presuppositions about the features of consciousness. We celebrate this aspect of his work. Our ordinary conception of consciousness is grounded in a specific evolutionary, developmental, and social history, mostly focused on typical human cases and a few familiar vertebrates, as commonly understood. Even if typical humans are normally fully and determinately phenomenally unified,<sup>25</sup> typical humans are a tiny corner of the architectural possibility space. What seems metaphysically necessary, as viewed from this corner, may prove instead to be a matter of contingent fact.<sup>26</sup>

---

<sup>25</sup> We take no stand on this issue here. Against unity even in typical human cases, see O’Brien & Opie 1998; Blackmore 2016.

<sup>26</sup> For helpful discussion, thanks to Tim Bayne, Ned Block, David Chalmers, Kenny Easwaran, Daniel Greco, P.D. Magnus, Kris Rhodes, Luke Roelofs, Jenelle Salisbury, and commenters on relevant social media posts on Facebook, Twitter, Bluesky, and The Splintered Mind.

## References

- Alter, Torin (2023). *The matter of consciousness*. Oxford University Press.
- Antony, Michael V. (2008). Are our concepts CONSCIOUS STATE and CONSCIOUS CREATURE vague? *Erkenntnis*, 68, 239-263.
- Baars, Bernard J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.
- Bayne, Tim (2010). *The unity of consciousness*. Oxford University Press.
- Bayne, Tim & David J. Chalmers (2003). What is the unity of consciousness? In A. Cleeremans, ed., *The unity of consciousness*. Oxford University Press.
- Blackmore, Susan (2016). Delusions of consciousness. *Journal of Consciousness Studies*, 23 (11-12), 52-64, 2016
- Bostrom, Nick (2006). Quantity of experience: Brain duplication and degrees of consciousness. *Minds and Machines*, 16, 185-200.
- Brook, Andrew & Paul Raymont (2001/2017). Unity of consciousness. *Stanford Encyclopedia of Philosophy*, Summer 2021 edition.
- Builes, David (2021). The world just is the way it is. *The Monist*, 104, 1-27.
- Carls-Diamante, Sidney (2017). The octopus and the unity of consciousness. *Biology and Philosophy* 32(6): 1269-1287.
- Carruthers, Peter & Rocco Gennaro (2001/2020). Higher-order theories of Consciousness. *Stanford Encyclopedia of Philosophy*, Fall 2023 edition.
- Chalmers, David J. (2002). Does conceivability entail possibility? In T.S. Gendler and J. Hawthorne, eds., *Conceivability and possibility*. Oxford University Press.
- Cochrane, Tom (2020). A case of shared consciousness. *Synthese*, 199, 1019-1037.

- Coleman, Sam (2013). The real combination problem: Panpsychism, micro-subjects, and emergence. *Erkenntnis*, 79, 19-44.
- Dainton, Barry (2000). *Stream of consciousness*. Routledge.
- Dainton, Barry (2014). Unity, synchrony, and subjects. In D.J. Bennett and C.S. Hill, eds., *Sensory integration and the unity of consciousness*. MIT Press.
- Dehaene, Stanislas (2014). *Consciousness and the brain*. Viking.
- Dennett, Daniel C. (1991). *Consciousness explained*. Little, Brown and Company.
- Dennett, Daniel C. (2005). *Sweet dreams*. MIT Press.
- Dennett, Daniel C. (2007). Heterophenomenology reconsidered. *Phenomenology and the Cognitive Sciences*, 6, 247–270.
- Dennett, Daniel C. (2016). Illusionism as the obvious default theory of consciousness. *Journal of Consciousness Studies*, 23 (11-12), 65-72.
- Dominus, Susan (2011). Could conjoined twins share a mind? *New York Times Magazine* (May 25): <https://www.nytimes.com/2011/05/29/magazine/could-conjoined-twins-share-a-mind.html>
- Fekete, Tomer, Cees Van Leeuwen & Shimon Edelman (2016). System, subsystem, hive: Boundary problems in computational theories of consciousness. *Frontiers in Psychology*, 7. DOI: <https://doi.org/10.3389/fpsyg.2016.01041>.
- Frankish, Keith (2016). Not disillusioned: Reply to commentators. *Journal of Consciousness Studies*, 23 (11-12), 256-289.
- Gendler, Tamar Szabó & John Hawthorne (2002). Introduction: Conceivability and possibility. In T.S. Gendler and J. Hawthorne, eds., *Conceivability and possibility*. Oxford University Press.

- Godfrey-Smith, Peter (2016). *Other minds*. Macmillan.
- Goff, Philip (2013). Orthodox property dualism + the Linguistic Theory of Vagueness = Panpsychism. In R. Brown, ed., *Consciousness inside and out*. Springer.
- Goff, Philip & Luke Roelofs (forthcoming). In defence of phenomenal sharing. In J. Bugnon et al. eds., *The phenomenology of self-awareness and conscious subjects*. Routledge.
- Heller, Mark (1996). Against metaphysical vagueness. *Philosophical Perspectives*, 10, 177-185.
- Hill, Christopher S. (2018). Unity of consciousness. *WIREs Cognitive Science*, 9 (5), e1465.
- Hirstein, William (2012). Mindmelding: Consciousness, Neuroscience, and the Mind's Privacy. *Oxford University Press*.
- Hurley, Susan L. (2003). Action, the unity of consciousness, and vehicle externalism. In A. Cleeremans, ed., *The unity of consciousness*. Oxford Academic.
- Jackson, Frank (1986). What Mary didn't know. *Journal of Philosophy*, 83, 291-295.
- Kang, Shao-Pu (2022). Shared consciousness and asymmetry. *Synthese* 200, (5), 1-17.
- Lockwood, Michael (1994). Issues of unity and objectivity. In C. Peacocke, ed., *Objectivity, simulation and unity of consciousness*. *Proceedings of the British Academy*, 83.
- Lycan, William G. How far is there a fact of the matter? *Journal of Consciousness Studies*, 29 (1-2), 160-169.
- Mallozzi, Antonella, Anand Vaidya, & Michael Wallner (2007/2023). The epistemology of modality. *Stanford Encyclopedia of Philosophy*, Summer 2024 edition.
- Masrour, Farid. The phenomenal unity of consciousness. In U. Kriegel, ed., *Oxford handbook of the philosophy of consciousness*. Oxford University Press.
- Nagel, Thomas (1971). Brain bisection and the unity of consciousness. *Synthese*, 22, 396-413.

- O'Brien, Gerard & Jonathan Opie (1998). The disunity of consciousness. *Australasian Journal of Philosophy*, 76, 378-95.
- Oizumi, Masafumi, Larissa Albantakis, & Giulio Tononi (2014). From the phenomenology to the mechanisms of consciousness: Integrated Information Theory 3.0. *PLoS Computational Biology*, 10 (5), e1003588.
- Parft, Derek (1984). *Reasons and persons*. Oxford University Press.
- Roelofs, Luke (2016). The unity of consciousness, within and between subjects. *Philosophical Studies*, 173, 3199-3221.
- Roelofs, Luke (2019). *Combining minds*. Oxford University Press.
- Roelofs, Luke (in draft). What's wrong with one-and-a-half minds?
- Rosenthal, David M. (2003). Unity of consciousness and the self. *Proceedings of the Aristotelian Society*, 103, 325-352.
- Salisbury, Jenelle Gloria (2023). *The unity of consciousness and the first-person perspective*. Doctoral Dissertation, University of Connecticut.
- Schechter, Elizabeth (2014). Partial unity of consciousness: A preliminary defense. In D.J. Bennett and C.S. Hill, eds., *Sensory integration and the unity of consciousness*. MIT Press.
- Schechter, Elizabeth (2018). *Self-consciousness and "split" brains*. Oxford University Press.
- Schwitzgebel, Eric (2007). No unchallengeable epistemic authority, of any sort, regarding our own conscious experience—contra Dennett? *Phenomenology and the Cognitive Sciences*, 6, 107-113.
- Schwitzgebel, Eric (2014). Tononi's Exclusion Postulate would make consciousness (nearly) irrelevant. Blog post at *The Splintered Mind* (Jul 16).

- Schwitzgebel, Eric (2015). If materialism is true, the United States is probably conscious, *Philosophical Studies*, 172, 1697-1721.
- Schwitzgebel, Eric (2023). Borderline consciousness, when it's neither determinately true nor determinately false that experience is present. *Philosophical Studies*, 180, 3415-3439.
- Schwitzgebel, Eric (2024). *The weirdness of the world*. Princeton University Press.
- Schwitzgebel, Eric, & Sophie R. Nelson (2023). Introspection in group minds, disunities of consciousness, and indiscrete persons. *Journal of Consciousness Studies*, 30 (9-10), 288-303.
- Searle, John R. (2000). Consciousness. *Annual Review of Neuroscience*, 23, 557-578.
- Simon, Jonathan A. (2017). Vagueness and zombies: Why 'phenomenally conscious' has no borderline cases. *Philosophical Studies*, 174, 2105–2123.
- Shani, Itay & Heath Williams (2022). The incoherence challenge for subject combination: an analytic assessment. *Inquiry*: <https://doi.org/10.1080/0020174X.2022.2124541>
- Sotala, Kaj & Harri Valpola (2012). Coalescing minds: Brain uploading-related group mind scenarios. *International Journal of Machine Consciousness*, 4, 293-312.
- Strawson, Galen (2003). What is the relation between an experience, the subject of experience, and the content of experience? *Philosophical Issues*, 13, 279-315.
- Tye, Michael (2003). *Consciousness and persons*. MIT Press.
- Vogel, Jonathan (2014). Counting minds and mental states. In D.J. Bennett and C.S. Hill, eds., *Sensory integration and the unity of consciousness*. MIT Press.
- Volz, Lukas J., & Gazzaniga, Michael S. (2017). Interaction in isolation: 50 years of insights from split-brain research. *Brain: A Journal of Neurology* 140(7): 2051-2060.
- Wundt, Wilhelm (1897/1897). *Outlines of psychology*, trans. C. H. Judd. Wilhelm Engelmann.

Yetter-Chappell, Helen (in draft). The view from everywhere: Realist idealism without God.