Why Consciousness Can't Overlap

Hedda Hassel Mørch, University of Inland Norway

DRAFT – under review

Last edited June 16, 2025

1 Introduction

Can consciousness overlap? That is, can a conscious system have parts that are also conscious, or itself be a part of a larger system that is?

This possibility is widely dismissed as highly counterintuitive (Putnam 1967; Block 1978; Merricks 2003; Unger 2004). But there have been relatively few attempts at backing the intuition up by argument. In philosophy and neuroscience, there are also a number of theories that directly challenge it, by suggesting or even implying overlapping consciousness, including functionalism (Block 1978), some accounts of split-brain cases, and various versions of panpsychism (Chalmers 2013; Goff 2017).

An exception comes from the Integrated Information Theory of consciousness (IIT) (Tononi et al. 2014; Albantakis et al. 2023). According to IIT's *Exclusion* postulate, "of any overlapping set of elements, only one set can be conscious" (Tononi et al. 2014: 3), specifically, the one with maximal integrated information. The postulate is defended by appeal to a mix of phenomenological observation and philosophical principles. However, many consider the defense unconvincing (Cerullo 2015; Schwitzgebel 2015; Bayne 2018; Blackmon 2021; Merker et al. 2021).

In this paper, I will offer a new defense of the claim that consciousness cannot overlap. I will first explain how various theories and arguments, across philosophy and neuroscience, seem to support overlapping consciousness. I will then examine arguments against overlap that have so far been offered and identify their main problems, with particular focus on IIT's defense of its *Exclusion* postulate.

I will then offer a new argument against overlap – building on Miller (2018) – which avoids previous problems. This argument is compatible with, and in many ways in the spirit of, IIT, and could therefore be incorporated into the theory's defense of *Exclusion*. But it does not rely on any premises specific to IIT, nor does it directly support any other postulate or part of the

theory. If the argument succeeds, the possibility of overlapping consciousness should therefore be universally rejected by any theory of consciousness.

2 Overlapping Consciousness

The question of overlapping consciousness is, more precisely, a question of whether *unified* consciousness can overlap. [1] Unified consciousness is when a variety of experiential qualities or contents – such as color qualities, sound qualities, thoughts, emotions and bodily feelings – are all subsumed under one and the same point of view,[2] or experienced together as part of a single experiential field – by one and the same *subject*, one might also add (how precisely the nature of subjects could be understood will be discussed in detail later).

Having unified consciousness can also be equated with the property of being conscious *as a whole*, as opposed to merely consisting of individually conscious parts, or merely being a part of larger conscious whole. The former is how we typically think of groups of people: the group consists of individually conscious people but the group as such is not (we typically think) conscious as a whole. The latter is how we typically think of, e.g., individual neurons within the conscious part of the brain (i.e., the NCC): these neurons are part of a larger conscious structure, but they are not individually conscious themselves.

The question of overlapping consciousness also concerns *synchronic* unity, i.e., the way all the experiences of a subject are unified at a given moment. It does not concern *diachronic* unity, understood as the way the experiences of a person at a time are (or are at least typically considered to be) unified with other experiences of the same person at other, later or earlier times; i.e., the relation (if any) constitutive of personal identity over time. Nor does it concern the kind of diachronic unity involved in the "specious present" (James 1890), i.e., in individual

[1] This is because overlap between entities with non-unified consciousness occurs more or less trivially. Any non-simple entity with unified consciousness will have parts (that it overlaps with) that are conscious in a non-unified way, and any two or more entities with unified consciousness will have a sum (that they overlap with) with non-unified consciousness, since non-unified consciousness simply consists in either being (a proper) part of an entity with unified consciousness or having (proper) parts that are. As we will see, overlap with between entities with unified consciousness is much less trivial. Also note that by "consciousness", I will mean phenomenal consciousness (or, in the case of theories that may – at least in some versions – be understood as eliminativist about phenomenal consciousness, such as functionalism, its closest functional equivalent).

[2] This is equivalent to what Bayne and Chalmers (2003) call *subsumptive* unity and Bayne (2010) calls *phenomenal* unity. Also note that this sort of unity could in principle subsume only one experiential quality or type of content (for example, we could conceive of a unified consciousness containing only one color quality and nothing else), i.e., it can but must not necessarily subsume a larger variety, but in typical (perhaps all actual human) cases, there will always be a variety.

moments of experience assuming these are not instantaneous but rather extended at least very briefly over time. [3]

To say that *synchronically unified* consciousness overlaps would be to say that there is a set of experiential qualities or contents, where (1) all of them are currently subsumed under one point of view/part of one experiential field and (2) a smaller subset of them is *also* currently subsumed under another point of view/part of another experiential field. Or, in terms of subjects, it is when some (but not all) of the experiences of one subject are also experienced by another subject, i.e., when two or more subjects share some (but not all) of the *numerically* same experiences (i.e., the same token experiences). For example, *Subject 1* might experience a color quality, a smell and an emotion, while *Subject 2* experiences the numerically same color and smell but not the emotion. Or alternatively, assuming these contents are reducible to further components, *Subject 2* may experience, e.g., some components of *Subject 1's* emotion but not others.[4]

3 Roads to Overlap

I will now consider the most important arguments and theories that support or imply overlapping consciousness (by which I mean overlap between synchronically unified consciousness, throughout the paper).

3.1 Functionalism

Functionalism is the view that a system is conscious if and only if it performs certain kinds of functions. It seems that most functions that could plausibly be associated with human consciousness, such as computational or psychological functions, can and often will overlap with other systems performing the same function. For example, according to Block's Chinese Nation argument (Block 1978), any function performed by the brain could in principle be

[3] This sort of unity is distinct from the diachronic unity involved in personal identity, which would be a unity holding between different experiential moments (most of which would be synchronically disunified with each other, and perhaps individually but not jointly specious). Still, the diachronic unity involved in the specious present (assuming moments are indeed temporally extended) is more closely related to synchronic unity than the kind involved in personal identity. The argument I will offer against overlapping synchronic unity may there also rule out overlapping diachronic unity of the sort involved in the specious present, but not the kind involved in personal identity (i.e., it may rule out overlapping specious presents, but not overlapping persons or subjects extended over time for longer than one specious present).

[4] It is also possible for *Subject 2* to experience additional qualities not experienced by *Subject 1* – one overlapping subject or its contents may be totally contained by the other, as in the examples given so far, but it need not be (they can also be related as a standard two-circle Venn diagram, as it were).

replicated by a vast number of individual people each performing the functional role of one neuron. If the group succeeded in doing this, it should, according to functionalism, gain consciousness, and this consciousness would overlap with the consciousness of the people in the group.

Block regarded the possibility of overlap, as illustrated by this scenario, as a "trouble" for functionalism, i.e., a reason to reject the view. Others have defended it as an acceptable implication (Schwitzgebel 2015; Schwitzgebel also gives additional arguments to support the implication).

3.2 The Integrated Information Theory – without *Exclusion*

The Integrated Information Theory (IIT) claims that a system is conscious if and only if it instantiates maximal integrated information, or maximal *phi* for short. Information, as IIT defines it, is very roughly a measure of internal determinism and non-degeneracy:[5] the extent to which the past and future states of a system are constrained by the internal state of the system itself as opposed to the external environment. Integration measures, again very roughly, the extent to which this internal determinism depend on the causal interconnectivity between the system's parts. A system has maximal integrated information if it has more integrated information than any overlapping system; in other words, if it has more integrated information than any of its own parts, as well as any larger system that it is itself a part of. [6]

If consciousness were to require mere integrated information, as opposed to a maximal level of it, IIT would imply that consciousness massively overlaps. This is because almost any system with mutually interacting parts will possess at least a small amount of *phi*, and IIT sets no threshold above which *phi* starts to matter for consciousness – any non-zero value will do. Without the maximality requirement, a small amount of consciousness would therefore have to be ascribed to almost any system. For example, not only would we have consciousness in some select area of the brain, but also in each of its subsystems, each of its neurons, each molecule within the neurons, and so on, all the way down to the microphysical level (Koch 2012: 132). It would also have to be ascribed to any larger area of the brain, the whole brain,

[5] Degeneracy is when multiple past states of a system could have produced its current state, whereas (full) non-degeneracy is when only one past state could have produced the current state. Non-degeneracy with respect to the past and internal determinism with respect to the future (i.e., when the current state can only produce one future state, or a few, if we are talking about a high but non-absolute level of it) both contribute to information in IIT's sense.

[6] See Mørch (2017) for a more elaborate non-technical introduction to IIT. See Albantakis et al. (2023) for the full details of the latest version of the theory.

the whole body, as well as larger encompassing systems such as society, the planet, the solar system, the galaxy and even the universe.

Given the maximality requirement, however, overlapping consciousness is not merely reduced but altogether ruled out, as among overlapping systems, there can only be one maximum of *phi* and hence one consciousness.[7] If some area of the brain has maximal *phi* (as IIT claims is clearly the case), then all the aforementioned overlapping systems, from neurons and particles to the larger organism, society or the galaxy, will be excluded from also being conscious (as a whole, or in a synchronically unified way, it bears repeating).

The maximality requirement is an integral part of IIT, and has never been questioned by the theory's authors, but several critics have argued that it is poorly supported and should be abandoned (Cerullo 2015; Schwitzgebel 2015; Bayne 2018; Blackmon 2021; Merker et al. 2021). So, whereas IIT itself does not support overlap (but rather strictly excludes it), arguments that the maximality requirement should be abandoned – while assuming or at least leaving open that the rest of the theory may be correct – do.[8]

3.3 Other Neuroscientific Theories – via the Problem of the Many

In philosophy, we find what Unger calls the problem of the many (1980), according to which most entities seem to overlap with an (often enormous) number of entities of the same sort as itself. Clouds, for example, have a large number of overlapping parts that, considered by themselves, would also qualify as clouds, such as what we typically think of as the cloud minus one vapor molecule, the same cloud minus another vapor molecule, minus two vapor molecules, and so on.

The same applies to neural or physical correlates of consciousness posited by most neuroscientific theories (Mørch 2019). The Global Workspace Theory (Baars 1993; Dehaene and Naccache 2001), for example, claims that consciousness correlates with the ignition of (or broadcasting of information within) a global workspace, a specific structure in which information can be distributed among different modules via a central hub. But almost any system that constitutes a global workspace will have a number of parts that considered in

[7] Except in the case of ties. According to the latest version of IIT, consciousness requires a *unique* maximum of *phi* (Albantakis et al. 2023). In cases of ties, consciousness gets attributed to the system with the next to highest *phi* relative to the tied systems.

[8] Many of the cited authors criticize other parts of the theory as well, but the criticism of the *Exclusion* postulate is still (at least somewhat) independent of criticism of other parts. Schwitzgebel and Blackmon, however, focus their criticism on the maximality criterion alone.

isolation would also constitute a global workspace, such as whatever global workspace structure we first consider minus one neuron, minus another neuron, minus two neurons, and so on. If instantiating a global workspace is sufficient for consciousness, all these overlapping systems should be conscious as well

Other theories of consciousness, such as the Recurrent Processing theory and the Orch OR-theory would face the same issue – unless they, like IIT, also posit a maximality requirement – according to which only the largest, or otherwise maximal among some other dimension, global workspace, recurrent process, orchestrated quantum collapse, or other proposed correlate would be conscious.[9] At present none of these theories posit such a requirement, hence overlap would seem directly implied by them.

3.4 Split-Brain Cases and Conjoined Twins

Split-brain patients are people who have had their corpus callosum, the set of fibers most directly connecting the left and right brain hemisphere, severed. In some experiments, these patients show behavior many have taken to indicate a splitting of consciousness, i.e., one unified consciousness residing in each hemisphere. For example, split-brain patients might verbally give one answer to a question, while their left hand points to an entirely different answer, indicating that the left hemisphere, which controls speech, does not have access to the same information as the right, which controls the left hand. In other respects, however, the hemispheres seem to share information. This has been interpreted by some as an indication that the two consciousnesses partially overlap (Nagel 1971; Lockwood 1989; Schechter 2014).

Also relevant are conjoined twins joined at the head, as well as the brain. The most famous case of this is the Hogan twins. These twins share a lot of sensory information, but not all; for example, one twin, Tatiana, is able to see the input of both eyes of the other twin, Krista, but Krista is able to see the input of one of Tatiana's eyes but not the other (Cochrane 2020). Several philosophers have taken this to indicate that the twins have partially overlapping consciousness (Montero 2016; Roelofs 2019; Cochrane 2020).

[9] An argument against introducing such a requirement would be that it renders consciousness non-intrinsic, since maximality (as well as being the largest of something, and so on) is an extrinsic property, i.e., one that depends on how things are outside the conscious system. Mørch (2019) shows that this could be avoided by adopting a non-reductionist view of consciousness, such as Russellian monism or dualism, according to which consciousness is not reducible to its non-intrinsic correlate. An option for those committed to reductionism (such as many proponents of many of the neuroscientific theories discussed here, presumably) is to simply accept that consciousness is not intrinsic.

3.5 Panpsychism and Cosmopsychism

In philosophy, panpsychist theories of consciousness have recently seen a revival, being proposed as a viable alternative to physicalism and dualism (Strawson 2006; Chalmers 2013). According to panpsychism, fundamental entities have fundamental forms of consciousness. The standard, sometimes referred to as "smallist", version of panpsychism assumes that microphysical entities such as particles are fundamental, and that particles therefore enjoy a simple and rudimentary form of consciousness. Complex macroconsciousness, of the sort we find in humans, is thought to arise from the combination of this simple microconsciousness, i.e., when microconscious particles are put together in the right way, such as in the brain.

This suggests that our own human consciousness would overlap with the microconsciousness belonging to the particles in our brain – unless these particles are posited to lose their individual, unified consciousness upon combination. Some theories of combination, such as the fusion view (Seager 2010, 2016; Mørch 2014, 2018), claim exactly this, but most other theories do not, and often explicitly acknowledge that overlap thereby does follow (Goff 2016; Roelofs 2019).

Another version of panpsychism is cosmopsychism (Shani 2015; Goff 2017), according to which the whole universe is fundamental and therefore enjoys a fundamental, but in this case highly complex, unified consciousness. Human and other non-cosmic forms of consciousness would thereby overlap with the cosmic consciousness – unless human and other non-cosmic forms of consciousness are taken to leave a "hole" in the cosmic consciousness (leaving it no longer strictly cosmic), but this is rarely if ever assumed by cosmopsychists.

4 Previous Arguments Against Overlap and Their Problems

Although overlap thereby enjoys fairly broad theoretical support, it is generally regarded as highly counterintuitive. This has been backed up experimental philosophy with respect to the general population (Knobe and Prinz 2008). Many philosophers have also dismissed it more or less out of hand simply based on its inherent implausibility (Putnam 1967; Block 1978; Unger 2004). But there have also been attempts at defending the intuition by more explicit arguments.

4.1 Moral Arguments

In philosophy, some have argued against overlap based on its allegedly implausible moral implications. Merricks (2003) makes the argument that if persons overlap, it follows that every

time you cut your hair you are actually committing murder – since there would be overlapping persons that consist partially of this hair – which is obviously absurd.

Overlapping consciousness is not quite the same as overlapping persons, especially since a person may be constituted also by stuff outside the brain (such as one's hair) whereas consciousness does not seem to be (*pace* some radical externalist theories). Simon (2018), however, makes a similar argument based directly on consciousness. He argues that, if consciousness could overlap, people with a higher number of neurons would contain more overlapping minds than people with fewer neurons, and therefore have higher moral status, which he also takes to be absurd.

These arguments will perhaps not be fully convincing to those otherwise open to the possibility of overlapping consciousness. Firstly, the absurdity of the alleged moral implications is also based on intuition, and the moral intuitions that speak against overlap indirectly are not clearly any stronger or well-founded than the original intuition directly against it. Secondly, it is not totally clear that these moral implications do follow. For example, Roelofs and Sebo (2024) argue that perhaps only token experiences (such as token pains) are morally relevant, i.e., they don't take on stronger moral relevance if a higher number of subjects have the experience, and overlap only multiplies the subjects of the same token experiences rather than experiences themselves.

4.2 The Argument from Phenomenal Holism

Another argument against overlap comes from Basile (2010). The argument takes as a premise that consciousness is holistic, in the sense that phenomenal qualities are determined or interpenetrated by their phenomenal context, i.e., any other phenomenal qualities they are experienced together with. For example, Basile claims, coffee in Naples will not taste exactly the same as coffee in Edinburgh, because in each case the taste would be embedded in a different overall experiential field subtly affecting the taste experience in different ways. He also assumes that the identity of a phenomenal quality is wholly determined by what it is like – i.e., that phenomenal qualities *are* essentially as they appear. If different phenomenal contexts will make a quality appear different, as per holism, it follows that they would also be different, such that the exact same quality cannot really occur in two different phenomenal contexts. But if a phenomenal quality is shared between two overlapping (but not exactly overlapping) subjects, it will occur in two different phenomenal contexts; hence such sharing would not be

possible. That is, a shared quality would appear and hence be different for each of the two subjects such that it cannot really be shared after all.

This is a powerful argument against overlap; however, the strong phenomenal holism it relies on does not obviously hold, and a number of doubts can be raised about it (Dainton 2010; Lee 2015; Cochrane 2020).[10]

4.3 IIT's Defense of *Exclusion*

A final notable argument against overlap can be found as part of the defense of IIT. Characteristically, IIT as a whole is supported by its authors not only by empirical evidence, but also by a philosophical or phenomenological argument. The argument is premised on a set of five *axioms* that describe features that, according to the argument, are necessary for any conscious experience, and are directly revealed as such by phenomenological observation or introspection. The axioms are then used to derive a set of *postulates*, describing physical features a system must possess in order to be conscious. The postulates are in turn mathematically formalized as the *phi* measure.

The maximality requirement follows from the postulate of *Exclusion,* according to which:

> Of all overlapping sets of elements, only one set can be conscious – the one … that is maximally irreducible … to independent components. [11] (Tononi et al. 2014: 3)

The postulate is derived from a corresponding *Exclusion* axiom, which has been formulated as follows:[12]

> Consciousness is exclusive: each experience excludes all others – at any given time there is only one experience having its full content, rather than a superposition of multiple partial experiences; each experience has definite borders – certain things can be experienced and others cannot. [13] (Tononi et al. 2014: 3)

[10] Dainton raises several doubts about the kind of phenomenal holism seemingly presupposed by Basile, but defends a different kind, which is, however, not clearly strong enough to support Basile's argument. See also Roelofs (2016) and Roelofs and Goff (forthcoming) for an argument that sharing is possible even given strong phenomenal holism, because a phenomenal quality experienced in one context by a Subject 1 and another by an overlapping Subject 2, the context experienced by Subject 1 may influence the experience of Subject 2 even though Subject 2 does not experience the context. Roelofs and Goff (forthcoming) also consider some additional philosophical arguments that can be or have been given against overlapping consciousness, and attempt to refute them, that I will not go into here.

[11] Maximal irreducibility is subsequently linked to maximal *phi.*

[12] This formulation is from IIT 3.0. Previous versions of IIT included a maximality requirement but were not as explicit about its source.

[13] Here I leave out a second part of the axiom that concerns the spatiotemporal grain of experience, which is less relevant for the overlap question.

The most recent version of IIT (Albantakis et al. 2023) puts the axiom more simply as: "Experience is *definite*: it is *this whole*."

In defense of the postulate, IIT has also invoked Occam's razor, i.e., made an appeal to parsimony:

> The exclusion postulate can be said to enforce Occam's razor (entities should not be multiplied beyond necessity): it is more parsimonious to postulate the existence of a single cause-effect structure over a system of elements – the one that is maximally irreducible from the system's intrinsic perspective – than a multitude of overlapping cause-effect structures whose existence would make no further difference. (Tononi and Koch 2015: 8)

Does the impossibility of overlap really follow from the axiom, Occam's razor, or their combination?

The *Exclusion* axiom claims, as we saw, that consciousness is definite or has definite borders. More specifically, it can be understood as presupposing that consciousness is unified (as is explicitly specified by the *Integration* axiom also part of IIT), and then adding that this unity is sharp and determinate as opposed to vague or indeterminate, such that there is always a definite fact of the matter as to whether some experience is included within this unity or not.

One might object to this claim that there are certainly cases of indeterminacy in experience,. For example, for objects or properties in our peripheral vision we are often not sure about whether we really experience them or not. Or, when perceiving a speckled hen (as in Chisholm's famous example (1942)), at a glance without counting its specks, we will experience an indeterminate number of specks.

IIT could respond to this in different ways. In some cases, it may appeal to the claim that not everything that is determinately present in a perceptual experience needs to be available to memory, conceptualization, further contemplation, and so on. In other words, phenomenal consciousness may "overflow" access consciousness, as IIT already assumes for independent reasons (Tononi et al. 2016), and in phenomenal consciousness everything is determinate even if not in access consciousness. In other cases, it could cast the phenomenon as a matter, not of it being indeterminate whether some determinate content is present in our experience, but rather of it being determinate that we are having an experience of some indeterminate content. That is, the indeterminacy is represented by experience; but the experience itself is still determinate.

In the speckled hen case, for example, we do not seem to be in a "superposition" of (as IIT calls it, i.e., in a state indeterminate between) an experience of, e.g., 150 specks, an experience

of 151 specks, 152 specks, and so on. We are rather determinately having an experience of an indeterminate number of specks (restricted, perhaps, to some range such as 100-200) – this would at least be a plausible candidate interpretation.

The most pressing problem for IIT's argument against overlap is therefore not about the axiom as such, but rather about how there clearly seems to be a gap between it and its corresponding postulate. The axiom's claim that consciousness has definite borders would fairly straightforwardly support a corresponding claim that physical systems that are conscious must also have definite borders, or roughly equivalently, that consciousness must be associated with a determinate physical property (i.e., one that does not admit of vague, borderline cases). But this is not what the postulate claims: the postulate claims that conscious systems cannot overlap with other conscious systems. And, as previously noted by Chalmers (2015), it does not follow from a system having definite borders that it cannot overlap with other systems.

To say that a system has definite borders is to say that it cannot have multiple alternative borders (that it would be "superposed" between), and thus that it cannot, in a sense, overlap with different versions of *itself*. It could still overlap, however, with *other* systems, such as one of its own parts. When a system has a part, the border around the whole will not thereby blur with the border around the part, or otherwise be rendered indefinite; rather, the whole and the part could each have their own, distinct and definite border.

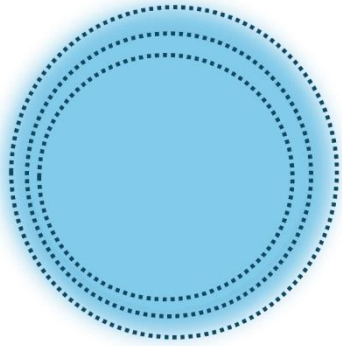

Figure 1: A system with an indefinite border around itself – this seems precluded by the Exclusion axiom

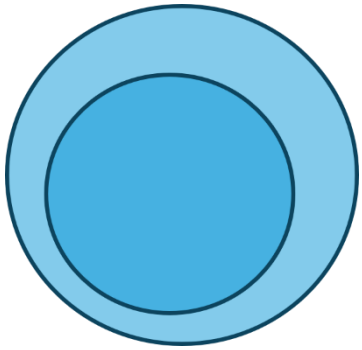

Figure 2: A system with definite border around itself, and another definite border around one of its parts – this is not precluded by the Exclusion axiom.

The *Exclusion* axiom thereby seems insufficient to support the postulate.

Could the additional appeal to Occam's razor, i.e., considerations of parsimony, help secure the link? This is doubtful as well. Positing multiple *distinct* entities for no good reason would clearly violate parsimony. But parts are precisely not distinct from the wholes they are parts of. Parts can also be regarded as necessary for the existence of the whole, or alternatively (on a holistic view) as necessitated by the whole, such that there *is* in fact good reason for positing parts given the whole. Overlapping parts of a system therefore do not seem excluded by Occam's razor.

5 A New Argument against Overlap

Opponents of overlapping consciousness are thereby in need of additional arguments to support their case. I will now offer an argument of this sort, that avoids the problems faced by the arguments discussed previously.

This argument, like IIT's, also starts from a purported "axiom", or truth about consciousness arguably discoverable from phenomenological introspection. Like IIT's *Exclusion* axiom, it also presupposes that consciousness is unified. It then adds the following, different, claim:

> No unified experience can have parts that are unified *in the same way* as the whole.

In other words, no part or subset of an experience can have a border, or "fence", around it, of the same type as the border around our entire experiential field.

Other kinds of unity can still characterize parts or subsets of an experience. For example, visual experiences are all unified, and clearly separable from, e.g., auditory or other sensory and non-

sensory experiences, but this is a *qualitative* sort of unity. A representation in perception can also be characterized by *objectual unity* (i.e., it can represent the features of an object, such as its color and its shape, as belonging to one and the same thing) (Bayne and Chalmers 2003). But the unity of an entire conscious field is a different kind of unity, as evidenced by the fact that this kind of unity typically obtains even though all of its contents are *not* unified in the qualitative, objectual or other sense that may hold for any of its parts. That is, my entire experiential field is currently neither qualitatively nor objectually unified, as it contains a number of very different qualities and representations of more than one object – but it is still unified in the sense of having a "border" around it separating my experiences from all other experiences I am *not* having (but are instead had by others, or nobody at all).

Unlike IIT's *Exclusion* axiom, this claim fairly straightforwardly supports a prohibition against overlap (i.e., something like the *Exclusion* postulate): If unified consciousness cannot have unified consciousness as parts, then physical systems that have unified consciousness should not be able to have parts that also have unified consciousness.

But is the claim true? It certainly *appears* true for all of my own experiences. But does this mean it is really true? And even if it happens to be true of all my (or all human) experiences, is it true of all possible experiences? I will offer an argument that the axiom is indeed universally true. The argument starts from distinguishing two different views about the nature of *subjects*.

It seems that all conscious experiences require a subject, in the sense of someone or something that *has* the experience or for whom/which the experience is *for*. But what are subjects, fundamentally? There are two main theories, or types of theories, about this.

First, there is *the substantive view* of subjects, typically associated with Descartes. On this view, a subject is an entity in its own right distinct from the experiences it is having. On one version of the view, this entity is understood as a container, that all the experiences of that subject are located within. On another version, it is understood simply a substance (or thing) that experiences inhere in.

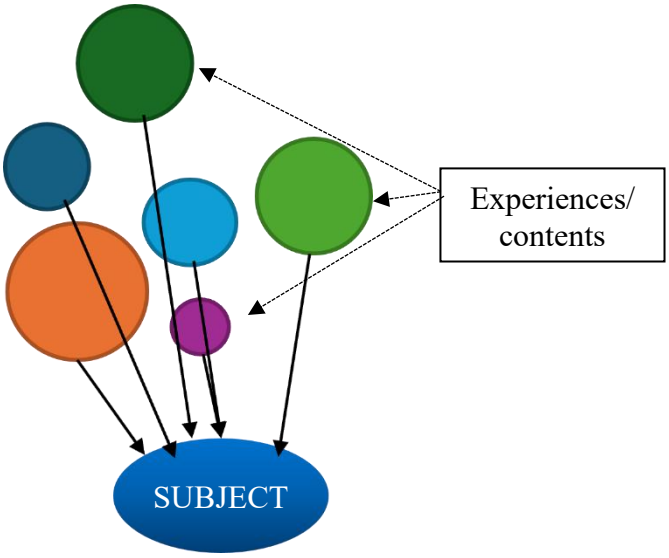

Figure 3: The substantive view, inherence version

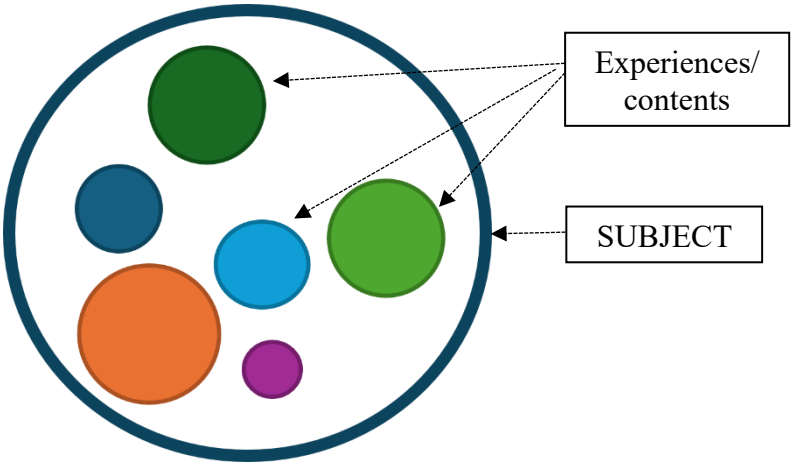

Figure 4: The substantive view, container version

Second, there is *the deflationary view,* typically associated with Hume. On this view, a subject is nothing more than a "bundle of perceptions", i.e., a set of experiences unified by relations amongst themselves, as opposed to by relations to some further entity (such as a container or substance). What kind of relation?[14] It seems that, in order to constitute a subject, a set of experiences must be, precisely, unified or co-conscious[15] with each other. A subject can thereby

[14] Some understand the bundle view as not requiring any type of special relation between the experiences. But this will lead to the standard objection that there must be something that "bundles the bundle", or something that makes all my experiences occur together without also including yours or anyone else's, which seems unanswerable without positing a relation that bundles experiences after all. In view of this, I will understand the bundle view as reducing subjects to a somehow "bundled" bundle rather than "unbundled" bundle, i.e., to experiences unified by some relations among themselves.

[15] Co-consciousness is typically defined as the relation of being experienced together from the same point view view (Dainton 2010), and can therefore regarded as identical with the unity relation.

be seen as reducible to a set of experiences related by a, perhaps primitive or fundamental,[16] unity or co-consciousness relation, whereas on the substantive view, unity should rather be reduced to the relation of inhering in or being contained by a primitive or fundamental distinct subject.

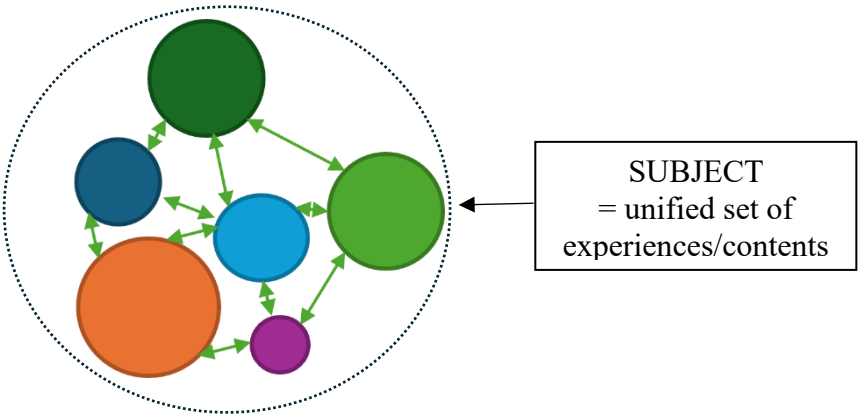

Figure 5: The deflationary view

Now, the substantive view of subjects seems fully compatible with overlap. Given this view, overlap would consist in a set of experiences inhering in or being contained by one subject, and a subset of the (numerically) same experiences also inhering in or being contained by another subject. There is nothing contradictory or otherwise obviously implausible in this idea. It also explains how it could be that, assuming our own consciousness overlaps with other smaller ones, this is not something we notice (i.e., why the claim proposed above, that no parts or subsets of our experiences are unified in the same way the whole, *appears* true): even if we have access to all the contents of own experience, we would not thereby have access to (in the sense of being able to detect the existence of) other subjects also experiencing some of the (numerically) same contents, because if subjects and their experiential contents are distinct, as implied by this view, access to the latter does not imply access to the former.

[16] One object that the definition of unity or co-consciousness as "being experienced together from a single point of view" presupposes the prior existence of a substantive subject as the locus of the point of view. But a unified point of must not obviously derive from a *distinct* and hence substantive subject. As Strawson (2008) argues, subjects could be regarded as identical to their experiences, or to a unified total experience, such that (in a paraphrase derived from Fodor (2007)), experiences simply have themselves.

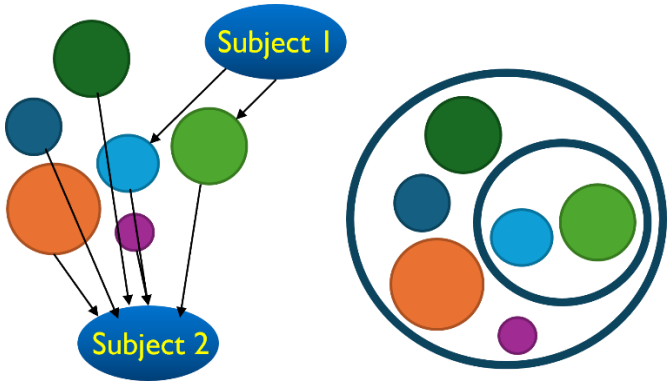

Figure 6: Overlapping subjects given the substantive view (inherence version on the left; container version on the right)

But the substantive view is arguably much less plausible than the deflationary view. The perhaps most important objection to the substantive view is the one originally posed by Hume, according to which we have no experience of any subject distinct from our experiences – the experiences (or experiential contents) themselves are all we have access to. A classic response to this – deriving from Kant – is that subjects must be posited in order to account for the unity of consciousness, which we do experience as holding between all the contents of our own experience, as well as not holding between our experiences and those of others. But this unity can also be accounted for by the deflationary view.

One might still posit substantive subjects even though we neither detect them or find a clear explanatory need for them. But doing so would then violate Occam's razor or parsimony. And even assuming that the substantive view can be justified in some way in and of itself, positing additional overlapping subjects, i.e., that some experiences inhere in or are contained by not just one but two (or more) distinct subjects, would still introduce additional lack of parsimony. More precisely, positing substantive subjects (rather than merely deflationary ones) is arguably *qualitatively* unparsimonious, because it introduces an entirely new fundamental kind of thing, whereas taking some experiences to have more than one substantive subject is, in addition, *quantitatively* unparsimonious, because it posits additional copies of a kind of thing that already exists. [17]

The deflationary view, on the other hand, does not seem compatible with overlap at all. To repeat, on the deflationary view, a subject consists in nothing more than a set of experiences (or experiential qualities or contents) related by a unity relation. But related how? It seems that

[17] Based on this, IIT's appeal to Occam's razor in defense of *Exclusion* could be interpreted as presupposing a substantive view of subjects. This would render it stronger, but it would still have the fatal weakness of ignoring the deflationary view.

the experiences belonging to one and the same subject must not only all be unified or co-conscious with each other, they must also *not* be unified or co-conscious with any other experiences. If some set of experiences, A, B and C, were all unified with each other but not with any other experiences, there is a clear sense in which there would be a subject identical with the bundle of A, B and C – a unified experience ABC. But if A, B, and C were also unified with a further experience D, the subject would actually extend to the entire ABCD bundle and not be limited to ABC only after all. In other words, if a set of experiences, such as ABC, are unified with each other but also with some additional experiences, such as D, then ABC should be considered merely part of larger subject ABCD rather than a complete subject in its own right.

It follows that the deflationary view should define a subject as a set of experiences that are (1) unified with each other, and (2) not unified with any other experiences.

However, as shown by Miller (2018),[18] this definition is not compatible with overlapping subjects. Given the above definition, overlap between a subject 1 and a smaller subject 2 would involve the following:

> **Subject 1** = a set of experiences that are (1) unified with each other and (2) not unified with any other experiences.
>
> **Subject 2** = a *subset* of the experiences of subject 1 that are (1) unified with each other and (2) not unified with any other experiences.

But this leads to a contradiction. It implies the existence of some subset of experiences, fully constitutive of subject 2 but only partially constitutive of subject 1, that are both unified with the other experiences of subject 1 (as follows from the description of subject 1) and *not* unified with them (as follows from the description of subject 2).

In other words, in order to be part of subject 1, that experiences contents A, B and C, the contents AB must be unified with C. But in order to constitute subject 2 (or for subject 2 to remain a subject in its own right instead of being absorbed into subject 1), the contents AB

[18] To be clear, what this paper contributes to the case against overlap beyond Miller's argument is to answer the objections below and to locate it within the larger argument that also takes into account the substantive view, as well as showing its relevance and applicability beyond the context of panpsychism and cosmopsychism, as Miller focuses on. Miller also concludes that the argument can be circumvented, i.e., that overlap may be possible despite of the argument, by relativizing unity to subjects (i.e., a set of experiences may be unified and exclusively bounded for one smaller subject while being unified and not exclusively bounded for another larger subject). This seems to presuppose that subjects are distinct from their experiences, or unified sets of them, i.e., the substantive view, and thus be vulnerable to the objections to accepting this view outlined above.

must *not* be unified with C. It follows that there is a subject ABC, then there cannot be an overlapping subject AB. This problem generalizes to any kind of overlap.

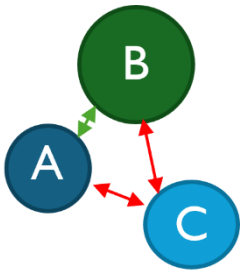

Figure 7: Overlap on the deflationary view. The existence of subject ABC requires the unity relations represented by the green and the red arrows. The existence of an overlapping subject AB requires the non-existence of the unity relations represented by the red arrows. The existence of both overlapping subjects is therefore impossible.

One might object that this contradiction can be resolved by positing different *kinds* of unity relations. If there are different kinds of unity relations, then subject 1 and its contents ABC could be related by unity relation type 1, while subject 2 and its contents AB can be related by unity relation type 2. This would involve no contradiction.

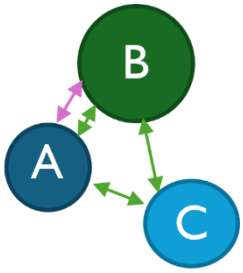

Figure 8: Overlapping subjects via two different types of unity relations. The green arrows represent unity relation 1 constitutive of subject 1. The purple arrow represents unity relation 2 constitutive of subject 2.

But as far as our own experience goes, there is only one kind of unity relation, the kind that seems to relate all our experiences but not extend to the experiences of others. When we imagine the consciousness of others, we imagine their experiences being related by the same kind of unity relation. We conceive of the unity of their consciousness as qualitatively identical to the unity of our own consciousness – the relata differ, but the relation remains the same. We cannot imagine or conceive of (at least not in any positive, qualitative detail) another kind of unity relation that is qualitatively different from our own and thereby be capable of relating

some smaller subset of our experiences at the same time as the total set of our experiences remains united by the standard or familiar unity relation.[19]

It therefore seems very implausible to posit such a relation, almost as implausible, one might argue, as positing distinct substantive subjects which we have not experienced either. In both cases, it seems they should not be posited simply to account for the possibility of overlap, a possibility for which there is no independent confirmation.

Another objection is that the deflationary view need not define subjects in terms of both the conditions specified above, but only in terms of the first one. That is, a subject can be defined as a set of experiences unified with each other, without requiring them *not* to be unified with any other experiences.

This view has been extensively defended by Roelofs (2019). A natural objection is that it has the counterintuitive implication that any set of experiences within our total experience (or other unified bundle) itself constitutes a subject. For example, my visual experiences would constitute a subject, as would my visual experiences minus one "pixel", my visual experiences minus two "pixels", half of my visual experiences, my hunger and one of my auditory experiences, and so on. But this may seem counterintuitive mainly in virtue of involving overlapping consciousness, so raising this as an objection in the current context could be dismissed as question-begging.

Roelofs also argues that although any set of unified experiences may constitute a subject in the minimal sense of being associated with a single point of view, they are not subjects in a number of other senses, such as constituting a coherent psychological structure or a person (this requires additional structural properties that do not follow from mere unity). This may contribute to making the overlap seem less counterintuitive.

But Roelof's view can still be subjected to the following additional objection. When we consider the original definition of a subject as a set of experiences (1) unified with each other but (2) not unified with any further experiences, the existence of a subject seems close to necessitated by both criteria of the definition being satisfied. When I conceive of a set of experiences related in accordance with both criteria, i.e., a unified and exclusively bounded (as

[19] If it were not qualitatively different, then the relation between A and B constitutive of the smaller subject would collapse, via the identity of indiscernibles, to the relation that already holds between A and B partially constitutive of the larger subject ABC. In terms of the illustration, if the purple arrow between A and B was not purple but green (i.e., not qualitatively different from the relation that also relates A and B to C), it would collapse into the green arrow, given that they are both qualitatively identical and have the same relata i.e., A and B.

one may put the requirement of not being further unified with anything else) bundle of experiences, then I am clearly conceiving of a subject in a minimal sense. It seems outright *inconceivable* that there should exist a unified and exclusively bounded set of experiences but no subject – in the sense of a single unified perspective, and thereby entity by which these experiences are all "had" at that moment.[20]

When we consider the alternative definition of a subject as a set of experiences that are simply (1) unified with each other, the existence of a subject is perhaps not ruled out by supposing there is something that satisfies this definition but *not* the original definition, i.e., a unified but *not* exclusively bounded bundle. But the existence of a subject does not seem anything close to necessitated. This is easy to see when considering our own experience. Any (proper) subset of my own experiences will constitute a unified bundle (i.e. satisfy the first requirement of the original definition and the only requirement of the alternative definition), as they will be unified with each other, but not an exclusively bounded one (i.e., not satisfy the second requirement of the original definition), since they will also be unified with everything else in my experience. I might be able to imagine that some such subset, such as my visual experience, constitutes a subject simply in virtue of constituting a unified (but non-exclusive) bundle, at least it does not seem to involve any clear contradiction. But I can also perfectly well conceive of it *not* constituting a subject.

The fact that a unified and exclusively bounded bundle is inconceivable without the existence of a subject, whereas a merely unified but non-exclusively bounded bundle is not, speaks strongly in favor of the former definition being the correct and sufficient one.

Summing up, then: The substantive view of subjects permits overlap, but this view is, firstly, unparsimonious in itself, and secondly, such that positing overlapping subjects, i.e., that some experiences have multiple subjects, would also be unparsimonious compared to positing only non-overlapping ones, i.e., that every experience has only one subject. On the deflationary view, in contrast, a subject is most plausibly characterized as a unified and exclusively bounded bundle, i.e., a set of experiences unified with each other but not with any other experiences. These kinds of bundles cannot overlap, on pain of contradiction.

[20] At least one the assumption that experiences are capable of "having themselves" (as discussed in footnote 15).

6 Implications for Neuroscience and Philosophy

This argument has consequences across neuroscience and philosophy. If the argument holds, it can firstly replace IIT's previous defense of *Exclusion.* It fits the axiomatic approach, by starting from a phenomenological observation of a seemingly necessary feature of our own experience – that no subsets of a unified experience are unified in the same way as the whole. It also incorporates its appeal to Occam's razor, but only invokes it against overlapping substantive subjects (as well as hypothetical additional unity relations beyond the familiar one). The new "axiom" also seems compatible with IIT's other axioms. By adopting this defense, IIT can respond to previous criticism according to which the *Exclusion* postulate should be dropped.

But even though the argument is compatible with IIT, it is in no way dependent on it, nor is it incompatible with any other theories of consciousness. As discussed, all other theories of consciousness in neuroscience that don't explicitly prohibit overlap will imply a massive amount of it, in view of the problem of the many. If the argument holds, this consequence is impermissible, so these theories must all be modified to avoid it. As already suggested, this could be accomplished by positing a maximality criterion analogous to *Exclusion,* according to which only the largest, most complex, or maximal (or perhaps minimal) along some other dimension, versions of their previously proposed correlates are conscious (or has unified consciousness, to be more specific).

Functionalism may have to do the same, i.e., posit that only maximal, or otherwise necessarily non-overlapping versions of particular kinds of functions constitute unified consciousness.

When it comes to split-brain cases, the argument rules out interpretations according to which these patients have partially overlapping consciousnesses (known as the partial unity model). The patients must either have two entirely distinct consciousnesses (known as the duality model), or their consciousness cannot really be split after all (known as the unity model). Both latter interpretations have been defended independently (see Schecter 2014 for an overview of defenses or endorsements of each model), so the fact that the argument rules out the interpretation according to which split-brain patients have split but overlapping consciousness should count against this interpretation rather than against the argument.

Conjoined twins such as the Hogan twins must be interpreted as at best having distinct token copies of parts of each other's experience (i.e., parts of their respective experiences are qualitatively but not numerically identical) rather than as sharing the same token experiences.

Even defenders of the view that the twins have partially overlapping consciousness have acknowledged that this alternative interpretation is consistent with what is known empirically about the case (Cochrane 2020; Roelofs and Sebo 2024).

In philosophy, the theory most impacted by this argument, besides functionalism, may be panpsychism. The argument implies that mental combination – the process by which, given standard panpsychism (i.e. setting cosmopsychism aside for the moment), simple microconsciousnesses come together to form a complex macroconsciousness like ours – cannot involve overlap; that is, the fundamental entities must lose their individual unified microconsciousness upon combination.

There are already theories of combination compatible with this, such as the aforementioned fusion view (Seager 2010, 2016; Mørch 2014, 2018). The fusion view is a version of emergent panpsychism, according to which mental combination involves strong emergence. Many panpsychists prefer constitutive views of combination, according to which mental combination involves only weak emergence. Most constitutive theories of combination imply that combined minds overlap with the fundamental minds that constitute it. But it seems possible for constitutive theories to posit criteria of combination that rule overlap, for example, by taking unified consciousness to be constituted by microphenomenal properties related in such a way that the overall structure instantiates some maximal property. The argument against overlap is therefore compatible with both emergent and constitutive standard panpsychism.

It is less compatible with cosmopsychism. Cosmopsychism is defined by the claim that the whole universe has a unified consciousness. If overlapping consciousness is impossible, it follows that our own unified minds cannot exist, since they would have to overlap with the cosmic mind. More or less approximate versions of cosmopsychism may still hold, according to which non-cosmic unified minds like ours create "holes" in the cosmic mind – "Swiss cheese" cosmopsychism, as one may call it. Another option is to think of non-cosmic unified minds as strongly emergent "outgrowths" of the cosmic mind, that are non-overlapping additions to rather than parts of it. On either view the cosmic mind is no longer fully cosmic as there would be minds that are part of the cosmos without being part of the otherwise cosmic mind.

This view might still satisfy some of the original motivations of cosmopsychism – mainly, that it is easier to see, according to its proponents, how minds like ours can derive from a larger or more complex mind than from a set of smaller, simpler minds, as per standard panpsychism;

i.e., that it arguably enables a better solution to the combination problem. Though there would be a limit to how many non-cosmic minds the universe can contain before the cosmic mind acquires so many "holes" or "outgrowths" that it no longer looks even approximately cosmic – for example, if not only living organisms, but also particles or other widespread physical structures, had unified consciousness.

7 Conclusion

Can consciousness overlap? Given the deflationary "bundle" view of subjects, overlap would involve a clear contradiction, as overlapping smaller subjects would have to be *not* unified with the larger subject in order to constitute its own subject, but also unified with the larger subject in order to be a part of it. This contradiction could in principle be resolved by positing multiple kinds of unity relations or by not requiring subjects to be *exclusively* bounded by the unity relation, but I have argued against both these moves.

Given the substantive view of subjects, overlap is possible, but this view is not only unparsimonious in itself compared to deflationary view; positing overlapping substantive subjects is also unparsimonious compared to positing only non-overlapping substantive subjects, such that even if the view is accepted, overlap should still be avoided.

If the argument holds, it vindicates IIT's *Exclusion* postulate, in a new and different way that is nevertheless compatible with and in the overall spirit of its original defense. It also implies that other theories in both neuroscience and philosophy should adopt a similar principle. Most of the affected theories should have no problem doing so, though with some exceptions, such as (strict) cosmopsychism.

References

Albantakis, Larissa, Leonardo Barbosa, Graham Findlay, Matteo Grasso, Andrew M. Haun, William Marshall, William G. P. Mayner, Alireza Zaeemzadeh, Melanie Boly, Bjørn E. Juel, Shuntaro Sasai, Keiko Fujii, Isaac David, Jeremiah Hendren, Jonathan P. Lang, and Giulio Tononi. 2023. Integrated Information Theory (Iit) 4.0: Formulating the Properties of Phenomenal Existence in Physical Terms. *PLOS Computational Biology* 19 (10).

Basile, Pierfrancesco. 2010. It Must Be True – but How Can It Be? Some Remarks on Panpsychism and Mental Composition. *Royal Institute of Philosophy Supplement* 85 (67).

Bayne, Tim. 2010. *The Unity of Consciousness*. Oxford University Press.

Bayne, Tim. 2018. On the Axiomatic Foundations of the Integrated Information Theory of Consciousness. *Neuroscience of Consciousness* 2018 (1).

Bayne, Tim, and David J. Chalmers. 2003. What Is the Unity of Consciousness? In *The Unity of Consciousness*, ed. Axel Cleeremans. Oxford University Press.

Blackmon, James C. 2021. Integrated Information Theory, Intrinsicality, and Overlapping Conscious Systems. *Journal of Consciousness Studies* 28 (11-12).

Block, Ned. 1978. Troubles with Functionalism. *Minnesota Studies in the Philosophy of Science* 9.

Baars, Bernard J. 1993. *A Cognitive Theory of Consciousness*. Cambridge University Press.

Cerullo, Michael A. 2015. The Problem with Phi: A Critique of Integrated Information Theory. *PLOS Computational Biology* 11 (9).

Chalmers, David J. 2013. Panpsychism and Panprotopsychism. *The Amherst Lecture in Philosophy* 8 (1-35).

Chalmers, David J. 2015. Integrated Information Theory: Some Philosophical Issues. In *Integrated Information Theory: Foundational issues*.

Chisholm, Rodrick. 1942. The Problem of the Speckled Hen. *Mind* 51 (204).

Cochrane, Tom. 2020. A Case of Shared Consciousness. *Synthese* 199 (1-2).

Dainton, Barry. 2010. Phenomenal Holism. *Royal Institute of Philosophy Supplement* 85 (67).

Dehaene, Stanislas, and Lionel Naccache. 2001. Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework. *Cognition* 79 (1-2).

Fodor, Jerry A. 2007. Headaches Have Themselves (Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism? By Galen Strawson Et Al). *London Review of Books*, 24 May 2007, 9-10.

Goff, Philip. 2016. The Phenomenal Bonding Solution to the Combination Problem. In *Panpsychism: Contemporary Perspectives*, eds. Godehard Brüntrup and Ludwig Jaskolla. Oxford University Press.

Goff, Philip. 2017. *Consciousness and Fundamental Reality*. Oxford University Press.

James, William. 1890. *The Principles of Psychology*. Vol. 1. Macmillan.

Knobe, Joshua, and Jesse J. Prinz. 2008. Intuitions About Consciousness: Experimental Studies. *Phenomenology and the Cognitive Sciences* 7 (1).

Koch, Cristof. 2012. *Consciousness: Confessions of a Romantic Reductionist*. MIT Press.

Lee, Geoffrey. 2015. Experiences and Their Parts. In *Sensory Integration and the Unity of Consciousness*, ed. Bennett Hill. MIT Press.

Lockwood, Michael. 1989. *Mind, Brain and the Quantum: The Compound "I"*. Vol. 99. Vol. 396. Blackwell.

Merker, B., K. Williford, and D. Rudrauf. 2021. The Integrated Information Theory of Consciousness: A Case of Mistaken Identity. *Behav Brain Sci* 45.

Merricks, Trenton. 2003. Maximality and Consciousness. *Philosophy and Phenomenological Research* 66 (1).
Miller, Gregory. 2018. Can Subjects Be Proper Parts of Subjects? The De-Combination Problem. *Ratio* 31 (2).
Montero, Barbara Gail. 2016. What Combination Problem? In *Panpsychism: Contemporary Perspectives*, eds. Godehard Bruntrup and Ludwig Jaskolla. Oxford University Press.
Mørch, Hedda Hassel. 2014. Panpsychism and Causation: A New Argument and a Solution to the Combination Problem (Doctoral Dissertation), Departement of Philosophy, Classics, History of Art and Ideas, University of Oslo, Oslo.
Mørch, Hedda Hassel. 2017. The Integrated Information Theory of Consciousness. *Philosophy Now* 121, https://philosophynow.org/issues/121/The_Integrated_Information_Theory_of_Consciousness.
Mørch, Hedda Hassel. 2018. Is the Integrated Information Theory of Consciousness Compatible with Russellian Panpsychism? *Erkenntnis*.
Mørch, Hedda Hassel. 2019. Is Consciousness Intrinsic?: A Problem for the Integrated Information Theory. *Journal of Consciousness Studies* 26 (1-2).
Nagel, Thomas. 1971. Brain Bisection and the Unity of Consciousness. *Synthese* 22 (3-4).
Putnam, Hilary. 1967. Psychological Predicates. In *Art, Mind, and Religion*, eds. W. H. Capitan and D. D. Merrill. University of Pittsburgh Press.
Roelofs, Luke. 2016. The Unity of Consciousness, within Subjects and between Subjects. *Philosophical Studies* 173 (12).
Roelofs, Luke. 2019. *Combining Minds: How to Think About Composite Subjectivity*. Oxford University Press.
Roelofs, Luke, and Philip Goff. forthcoming. In Defence of Phenomenal Sharing. In *The Phenomenology of Self-Awareness and the Nature of Conscious Subjects*, eds. Martine Nida-Rümelin and Julien Bugnon. Routledge.
Roelofs, Luke, and Jeff Sebo. 2024. Overlapping Minds and the Hedonic Calculus. *Philosophical Studies* 181 (6).
Schechter, Elizabeth. 2014. Partial Unity of Consciousness: A Preliminary Defense. In *Sensory Integration and the Unity of Consciousness*, eds. David Bennett and Christopher Hill. The MIT Press.
Schwitzgebel, Eric. 2015. If Materialism Is True, the United States Is Probably Conscious. *Philosophical Studies* 172 (7).
Seager, William. 2010. Panpsychism, Aggregation and Combinatorial Infusion. *Mind and Matter* 8 (2).
Seager, William. 2016. Panpsychist Infusion. In *Panpsychism: Contemporary Perspectives*, eds. Godehard Brüntrup and Ludwig Jaskolla. Oxford University Press.
Shani, Itay. 2015. Cosmopsychism: A Holistic Approach to the Metaphysics of Experience. *Philosophical Papers* 44 (3).
Simon, Jonathan A. 2018. The Hard Problem of the Many. *Philosophical Perspectives*.
Strawson, Galen. 2006. Realistic Monism: Why Physicalism Entails Panpsychism. *Journal of Consciousness Studies* 13 (10-11).
Strawson, Galen. 2008. What Is the Relation between an Experience, the Subject of the Experience, and the Content of the Experience? In *Real Materialism*. Clarendon Press.
Tononi, Giulio, Larissa Albantakis, and Masafumi Oizumi. 2014. From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLOS Computational Biology* 10 (5).

Tononi, Giulio, Melanie Boly, Marcello Massimini, and Christof Koch. 2016. Integrated Information Theory: From Consciousness to Its Physical Substrate. *Nature Reviews Neuroscience*.

Tononi, Giulio, and Christof Koch. 2015. Consciousness: Here, There and Everywhere? *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 370 (1668).

Unger, Peter. 1980. The Problem of the Many. *Midwest Studies in Philosophy* 5 (1).

Unger, Peter. 2004. The Mental Problems of the Many. In *Oxford Studies in Metaphysics, Vol. 1*, ed. Dean Zimmerman. Oxford: Clarendon Press.