


This article was downloaded by: 10.3.97.143

On: 06 Dec 2023

Access details: *subscription number*

Publisher: *Routledge*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK



## **The Routledge Handbook Of Consciousness**

Rocco J. Gennaro

### **The Unity of Consciousness**

Publication details

<https://www.routledgehandbooks.com/doi/10.4324/9781315676982-28>

Elizabeth Schechter

**Published online on: 26 Mar 2018**

**How to cite :-** Elizabeth Schechter. 26 Mar 2018, *The Unity of Consciousness from: The Routledge Handbook Of Consciousness* Routledge

Accessed on: 06 Dec 2023

<https://www.routledgehandbooks.com/doi/10.4324/9781315676982-28>

**PLEASE SCROLL DOWN FOR DOCUMENT**

Full terms and conditions of use: <https://www.routledgehandbooks.com/legal-notices/terms>

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

# THE UNITY OF CONSCIOUSNESS

*Elizabeth Schechter*

## 1 Introduction

I turn my head to gaze out the screen door to my left. It's a pretty scene—grass, rocks, flowers, and woods beyond—but I'm not inspecting it closely. Out of the corner of my eye I see someone crossing the kitchen, and hear someone else walking on the floor above me. Water is running somewhere. I've paused to consider whether I should go to the supermarket before or after lunch. Knowing B\_\_\_, she will be happy either way. Outside, my dog runs by in a happy white flash, and I realize I've been sitting with one leg bent under me. It's starting to tingle. Shifting position, I decide to get the grocery shopping done with.

Consider this *5-second episode* a window into my conscious experience, encompassing a number of elements. Some are perceptual: I hear someone walking overhead; I become proprioceptively aware of my posture. Some are cognitive: I weigh considerations, make a prediction about someone else's preferences. Some are metacognitive: I recognize that I'm looking outside only absentmindedly. Some elements are agential: I make a decision. Some seem motoric: I'm conscious of repositioning myself in my chair.

We can consider these and other elements of my experience individually, or we can consider the whole of my experience over the course of or at any moment within the episode. *Experiential wholes* are indeed not homogeneous or simple, but incorporate a multitude of *experiential elements*.

Questions about the *unity of consciousness* concern relationships between the individual elements comprised by an experiential whole, between those elements and that whole, and between these experiential phenomena and their experiencing subjects.

These questions can be sorted into six rough categories. First is the *taxonomy question*: if we say that I enjoyed a unified consciousness throughout the 5-second window, what feature or quality of my experience do we mean to indicate? The *analysis question* concerns one particular kind of conscious unity—phenomenal unity—and asks which non-phenomenal relation it is equivalent to. There are *metaphysical questions*, focusing especially on the relation of conscious unity relations. *Mechanism questions* concern the functional and neural bases of conscious unity. Whether and how consciousness can fail to be unified is the *disunity question*. *Identity questions* concern the relationship between conscious unity and subjects of conscious experience.

## 2 Taxonomy Question: What We Talk about When We Talk about Unity

Conscious experience has a number of unified and unifying features (Bennett and Hill 2014; Brook and Raymont 2014). Start with the fact that we seem to perceive things as standing in ordered *spatial* relations to each other and to ourselves. Or consider that for any two concurrent elements of one's experiential whole, it seems that one can demonstratively co-refer to them ("I am experiencing *those*"). For that matter, we don't experience the world as a mere jumble of features and occurrences, but as containing coherent objects participating in coherent events. Something seems to promote *mutual consistency* among the various elements of a single experiential whole, in such a way that the experience affords a coherent perspective on the world. Consider *binocular rivalry*. If two images, each of a different object, are presented simultaneously to a subject, one to each eye, the subject is not thereby induced to undergo a single fantastical visual experience of two material objects occupying the very same region of space-time. Typically, instead, subjects spontaneously alternate between seeing first one object in the region, and then the other (see Logothetis 1998 for review). Finally one might argue that the unity of consciousness is essentially tied up with agential and rational unity more broadly (Shoemaker 1996).

On the other hand, it seems to many philosophers that even if your consciousness failed to be unified in any or all of the foregoing senses, it might remain unified in just this one respect: that there would still be *something that it was like* (Nagel 1974)—some *one* thing—to be you experiencing whatever you're experiencing.

Conscious unity in this last sense is called *phenomenal unity*. Just as philosophical work on consciousness most often concerns phenomenal consciousness, philosophical work on the unity of consciousness most often concerns phenomenal unity. This is no coincidence. Phenomenality is puzzling, and philosophers love puzzles.

Phenomenal properties are famously difficult to define (see Block 1978: 281; cf. Dennett 1988), other than by ostension, that is, by pointing to examples. Consider this contrastive pair of cases (modified from Fodor 1981). Suppose you open your eyes in a hotel room one morning and in the first moment of consciousness, before you have any thoughts at all, you visually experience the sparsely furnished room, whose walls are all painted *red*. When you return to your hotel that evening you're informed that there was a problem with your room and that your things have all been moved to a new room with the same layout and furnishings and dimensions. The next morning you open your eyes and in the first moment of consciousness, before you have any thoughts at all, you visually experience the sparsely furnished room, whose walls are all painted *green*. Suppose that your old room and your new room are otherwise visually indistinguishable, and consider only your first moment of consciousness, before you think, "What is this unfamiliar *red* room?" on the first morning and "What is this unfamiliar *green* room?" on the second morning. The subjective character of your experience on the first morning would presumably be different from the subjective character of your experience on the second morning. This difference in subjective character is, ostensibly, a *phenomenal* difference, a difference between the phenomenal properties of your experience on those two days.

I will take phenomenal unity to be a *phenomenal* relation between phenomenal properties, in the sense that the phenomenal unity of two elements of experience makes a phenomenal difference to their subject. What makes it difficult to characterize or define phenomenal unity is that the phenomenal difference it makes is so abstract: presumably any two phenomenal properties can have the further property of being phenomenally unified, despite incredible diversity among phenomenal properties and thus among the phenomenal characters of their unified pairs.

This makes it especially tempting to try to articulate what phenomenal unity is by pointing to a pair of contrastive examples, as I did for phenomenal properties above. The most helpful pair of contrastive examples would be a case in which two elements of experience were phenomenally unified and a case in which they were not. The problem is that for various reasons it is not clear that the elements of a subject's experience at a time could ever fail to be phenomenally unified, nor that we can even imagine a case in which they so failed (see Sections 6 and 7).

To characterize phenomenal unity, then, reference is instead often made to there being something that it's like for the subject of a unified consciousness to experience *together* everything they experience, e.g.: "Experiences, when they occur simultaneously, do not occur as phenomenal atoms but have a conjoint phenomenology—there is something it is like to have them *together*" (Bayne 2010: 280; emphasis added). During the five-second episode, say, there wasn't just something it was like for me to see someone walk across the kitchen and something it was like for me to hear someone walking overhead. There was something it was like for me to see someone walk across the kitchen *while* hearing someone walking overhead.

Unfortunately it is very difficult to provide an elaborated description of this experienced "togetherness." On pain of redundancy, it can't just mean that the elements are experienced simultaneously or by the same subject. We might say that if you have a unified consciousness, then everything you experience is experienced as part of a *unified whole*. But a whole *what?* And again—unified *how?*

### 3 The Analysis Question: What Does Phenomenal Unity Reduce to or Consist In?

The ineffability of phenomenal unity makes it desirable that it should consist in some non-phenomenal relation. The most basic proposal is that phenomenal unity can be understood in terms of the conjunction of the contents of experience, a proposal explored by, among others, Hurley: "if and only if a conscious state with content *p* and a conscious state with content *q* are co-conscious [i.e. unified], then there is a conscious state with content *p and q*" (Hurley 1998: 117; citing Chisolm 1981; see also Tye 2003). This analysis of phenomenal unity is presumably entailed by representationalist accounts of consciousness, whose proponents believe that phenomenal properties reduce to causal-intentional properties.

Another proposal is that phenomenal unity might reduce to or consist in something like *feature binding* (or similar—see Revonsuo 1999), the process whereby multiple features of objects come to be accurately associated with each other in perception. Suppose I am watching my friendly white dog trot toward my wary gray rabbit. To experience this, I must perceive their colors, shapes, and movements—and as it happens, color, shape, and motion are all represented in different regions of the brain and in partial independence of each other; one can lose the ability to perceive one of these features without losing the ability to perceive the others, for instance (e.g. Zihl, Cramon, and Mai 1983). The so-called *binding problem* concerns how the features we experience come to be correctly bound, or associated, in perception—so that I perceive my *rabbit* and not my dog as *gray*, my *dog* and not my rabbit as *trotting*—indeed, so that I perceive colored moving objects at all, as opposed to a mere jumble of color, shape, motion. So, again, one proposal is that phenomenal unity consists in this kind of representational integration.

The philosophical literature contains several other potential analyses of the phenomenal unity relation. Nagel referred to our assumption that "for elements of experience...occurring simultaneously or in close temporal proximity, the mind which is their subject can also experience the simpler relations between them if it attends to the matter" (Nagel 1971: 407; emphasis omitted; and see again Tye 2003). We could pull from this the proposal that the phenomenal

unity of elements of experience consists in their *access unity*, which is an extension of the concept of *access consciousness* (Block 1995). Very roughly, a mental state is said to be access conscious when its subject can introspect or report it and use it in reasoning and the rational guidance of action. Access conscious states can be contrasted with the multitude of mental states we know about only from experimental or theoretical work demonstrating that they play roles in the non-rational guidance of our behavior and automatic responses, or in determining what we consciously experience. So, two experiential elements are *access unified* if they (or their contents) can *jointly* figure in reasoning. Perhaps this is what phenomenal unity amounts to.

Marks meanwhile suggested that two experiences “belong to the same unified consciousness only if they are [or could be] known, by introspection, to be simultaneous” (Marks 1981: 13). We might pull from this an analysis of phenomenal unity into *awareness unity*, an extension of the concept of *higher-order awareness* (Rosenthal 1986). According to higher-order theories of consciousness, a mental state is conscious in virtue of its (occurrent or dispositional) relation to a second mental state which is about it, or which takes it as its content. So, two elements of experience are *awareness unified* if they are or are disposed to be made the common object of a single higher-order mental state.

Although each of these candidate analyses of phenomenal unity picks out an interesting respect in which human consciousness is unified, none is widely agreed to offer an analysis of phenomenal unity. The trouble is that some philosophers believe that phenomenal unity could conceivably exist in the absence of any other kind of conscious unity: access unity, awareness unity, feature-binding, and perhaps even the conjunction of contents (see discussion in Bayne 2010). Indeed, philosophers have sometimes concluded that the phenomenal unity of experiential elements admits of no analysis—that it is a basic relation (Dainton 2000). Although most philosophers seem reluctant to accept this, the literature on the analysis question consists largely of arguments rejecting various candidates: not spatial or representational unity, not object or introspective unity, not access or awareness unity. (For taxonomies and critical discussions, see Bayne and Chalmers 2003; Tye 2003.)

The central difficulty concerns the notion of the phenomenal generally. It may not be essential to the concept of phenomenal unity that it should have any particular causal profile. Of course, this might prompt suspicion about the coherence of the concept of phenomenal unity (just as there has been about the concept of the phenomenal, generally; e.g. Dennett 1988, Church 1995).

Debates about phenomenal unity meanwhile inherit all of the controversies about phenomenal consciousness generally. For instance, suppose that there is no such thing as *cognitive phenomenology*: that coming to believe something, or judging or doubting something, or making a decision, have no intrinsic phenomenal characters—unlike, say, conscious seeing or conscious hunger. If this is right, then an account of phenomenal unity can simply ignore the *propositional attitudes* (beliefs, judgments, doubts, decisions—mental states that can be rational or irrational). This in turn makes it more plausible that phenomenal unity should reduce to a kind of *spatial unity* of (the contents of) experience. This analysis is less plausible where propositional attitudes are concerned, however. Thus, uncertainty about the contents of phenomenal consciousness affects the plausibility of competing accounts of phenomenal unity.

#### 4 Metaphysical Questions: Conscious Unity Relations at a Time and over Time

Every experience soon fades from consciousness. Most are never recalled. This makes it unlikely that all of one’s experiences stand in any relation to each other than mere succession. On the other hand, it seems possible that all of one’s experiences at a given moment are substantively

unified in some respect. It is thus standard for philosophers to distinguish between conscious unity *at a given time*, or synchronically, and conscious unity *over time*, or diachronically. Many metaphysical questions about conscious unity concern these temporal concepts.

Conscious experience has both objective and subjective temporal features. A given experience occurs at a particular time and has a particular duration, and these are objective temporal properties. Some experiences meanwhile seem to their subject to be momentary or to drag on, to come before or after others, and these are subjective temporal properties. Naturally objective and subjective temporal features are connected but not identical. For instance, consciousness takes time, so an event that I experience as happening now, in the *subjective present*, may in fact have happened in the very recent *objective past*.

Discussions of conscious unity relations between experiences *at a moment in time* often don't specify whether the moment referred to is one of objective or subjective time. That is, is synchronic conscious unity a relation between all the objectively simultaneous experiences of a subject? Or between all the happenings that the subject's experience presents to him as occurring simultaneously?

Many philosophers believe that we don't experience the present moment as strictly instantaneous but rather as a very brief interval; the term "specious present" is used to refer to the subjective present when it is conceptualized in this way. (The term is associated with William James 1890, though contemporary usages don't line up neatly with his.) It is tempting to conceptualize the subjective present as an interval in order to account for the difference between *successive experiences* and *experiences of succession*. Imagine watching someone dance under a strobe light: you have first a visual experience of the dancer in one position, then a visual experience of darkness, then a visual experience of the dancer in a new position. Your extended experience, of repeatedly seeing the dancer *having moved*, will be different from an experience of seeing a *dancer moving*, in ordinary light. Or imagine hearing a piano key hit repeatedly, versus hearing a key pressed and held. It is tempting to cast these differences in temporal terms. In the first kind of case, you perceive change (or stasis) across intervals of experiencing, but not within any one interval. In the second kind of case, you perceive change (or stasis) within an interval.

A philosopher employing the concept of the specious present for this purpose may view himself as offering an account of very short-term diachronic unity (Dainton 2000). Indeed, within the momentary interval of the specious present, phenomenal unity might even be *transitive*, such that two elements of experience that are both unified with a third must also be unified with each other. But if phenomenal unity were transitive throughout any arbitrary interval of time, then what it's like for, say, me to blow out candles on my 40th birthday cake, necessarily depends upon whether or not I experienced blowing out candles on my 4th birthday cake. This seems implausible. It may then be that for a subject to have a phenomenally unified consciousness over any extended period of time is simply for him continuously to have a consciousness that is phenomenally unified within the specious present (Dainton 2000; Dainton and Bayne 2005).

Metaphysical questions about conscious unity often concern the relation of conscious unity relations. The literature on conscious unity often takes it to be a relation between *experiences*, but this is not without controversy, since there is no universal agreement about the identities of experiences. According to the "many-in-one" position, my experiential whole is an experience, and so are the numerous elements it contains (Bayne 2010; Lockwood 1989). According to the "only-one" position, only the whole of my experience is itself an experience; its elements are not (Tye 2003). Finally, there is a "many-only" position, according to which I have a multitude of experiences at every moment that are unified without being somehow incorporated into a single overarching experience.

This disagreement concerns the individuation of experiences. Philosophical works on conscious unity mainly individuate experiences using a “tripartite” conception (Bayne 2010) that identifies them as particular phenomenal characters of particular subjects at particular times (Dainton 2000). That is, given a subject and a time, the tripartite conception individuates experiences on the basis of their phenomenal characters. This rules out certain possibilities, like the possibility that a single subject might have at a given moment in time multiple experiences with the same character (cf. Schechter 2013). The tripartite conception does not yield a single way of identifying experiences, however.

One might suppose that the tripartite conception rules out the only-one view: how could my experience of my white dog be the very same experience as my experience of my gray rabbit, given the phenomenal difference between experienced whiteness and experienced grayness? But it begs the question against the only-one view to read “experienced whiteness” as “experienced whiteness *alone*, in the absence of any experienced grayness.” According to the only-one view, an experience of my dog can equally be an experience of my rabbit, in the same way that my photograph of my dog can equally be a photograph of my rabbit—by being a photograph of them both! In itself, the tripartite conception of experiences is consistent with this position.

It is also consistent with the many-only view. It is true that my experiential whole, at any moment, has a phenomenal character, *W*, that differs from the phenomenal character of any one of my experiences, and also from the mere (disunified) sum of phenomenal characters associated with each of those experiences. The many-only view accepts this, however: it simply denies that *W* is the character of an experience. According to the many-only view, *W* is a *unity* of the phenomenal characters of my multitude of experiences, without itself being an experience.

Naturally, the tripartite account is also consistent with the many-in-one view. In fact, one intuitive metaphysics of the phenomenal unity relation casts it as a mereological relation between experiences. Bayne has offered this kind of account: two experiences are unified in virtue of being *subsumed* by another experience, with an experiential whole subsuming all (other) experiences of its subject (Bayne 2010).

The differences between these accounts must be subtle. Even the only-one view allows multiplicity somewhere, if only in the multiplicity of things that I experience in the world. Even the many-only view allows that a subject’s consciousness is unitary in the sense that she has, say, only a single perspective or stream of consciousness. It is therefore unclear what hinges upon whether we identify as *experiences* experiential wholes, experiential elements, or both. Perhaps some objective basis for individuating experiences could be drawn from empirical generalizations about, say, working memory capacity. This kind of possibility has not been much explored by philosophers of conscious unity, however, who may believe that since experience is essentially subjective, a subject should be able to tell precisely on the basis of her own experience just how many experiences she has (but see Hurley 1998; cf. Bayne 2010).

## 5 The Mechanism Question: How to Explain Conscious Unity

Further questions about conscious unity concern its broadly causal explanation, including its neural basis. Answers depend upon the kind of conscious unity at issue. Suppose that we understand conscious unity to be the product of feature binding. Then an answer to the *mechanism question* consists of an articulation of the mechanism of feature binding: say, two elements are unified when spatial attention selects both of them as concerning goings-on at a single location in space, thus “tagging” their contents, for further processing, as being features of a common object (Treisman 1998; Treisman and Gelade 1980). If one instead takes conscious unity to be a

kind of access unity, then two elements may be unified when their contents are jointly broadcast by attention into the “global workspace” (Baars 1988), or when they are simultaneously encoded in working memory. If conscious unity is taken to be a kind of unified awareness of experiences, then two experiential elements may be unified when their contents become the object of a single higher-order state.

On the other hand, if conscious unity is taken to be phenomenal unity, and if phenomenal unity is not analyzable into any non-phenomenal relation, then it's not immediately clear how to explain it. Here again there exists a parallel between phenomenal unity and what's often said to be true of phenomenality generally: that any causal accounting of it may leave an *explanatory gap* (Levine 1983).

One preliminary mechanism question about phenomenal unity is whether it even requires explanation in addition to whatever causal accounting is offered for phenomenal consciousness generally, or whether, instead, the mechanism of phenomenal unity is just the mechanism of phenomenal consciousness. Bayne (2010) uses the terms “atomism” and “holism” to refer to the former and to the latter possibility, respectively.

Holism is a more sensible proposal with respect to synchronic conscious unity. Any kind of extended diachronic conscious unity will presumably require mnemonic integration not necessary for experience itself. On the other hand, it seems possible that the mechanism of consciousness is also the source of its unity at a given moment.

Neither holism nor atomism need be true of *all* kinds of conscious unification. Consider *multisensory integration*—a kind of perceptual binding—as it occurs when, say, one experiences an audiovisual recording designed to induce the so-called *McGurk effect* (McGurk and MacDonald 1976). In one such video, a sound reel of an actor speaking the syllables *ba-ba* plays over a visual recording of him mouthing the syllables *ga-ga*. To the perceiver, however, the man seems to be saying *da-da*. The perceiver does not hear one thing, see another, and then struggle to work out what the speaker is actually saying. Rather, what the speaker is saying is worked out by the perceiver's perceptual systems, prior to conscious experience.

Multisensory integration and perceptual binding more generally have sometimes been proposed to integrate what would otherwise be disunified conscious contents (Bartels and Zeki 1998), and they do clearly contribute to the *coherence unity* of consciousness, creating for the experiencing subject a coherent world of objects and events. Thus atomism could be true for coherence unity. On the other hand these mechanisms of coherence unity are not well suited to explain phenomenal unity. Multisensory integration means that in the perceiver's actual experience, there is not a *ga-ga* thing that it is like to *see* the speaker speak unified with a *ba-ba* thing that it is like to *hear* the speaker speak. In the perceiver's experience, there is just the *da-da* thing. Phenomenal unity relates *already experiential* elements: what it's like to see the actor in the video say *da-da* and what it's like to hear him speak *da-da*.

Since phenomenal unity is a relation between the already-conscious, philosophers of phenomenal unity seem inclined toward holism. Among other things, if atomism were true, then one would expect disorders or impairments characterized by having a multitude of elements of experience not incorporated into any experiential whole, and there is no evidence that such cases exist (see Section 6 on the “disunity question”).

The debate between atomism and holism would be more readily answered with an agreed-upon theory of consciousness at hand. There has been significant convergence around particular neurofunctional theories of *access* consciousness, perhaps especially the global neuronal workspace model of Dehaene and colleagues (e.g. Dehaene and Naccache 2001), developed out of Baars's “global workspace” model (Baars 1988). Some philosophers believe that access consciousness is the only kind of consciousness that exists (e.g. Dennett 2001), or at least



that the machinery of access consciousness partially explains phenomenal consciousness (e.g. Carruthers 2017), though other philosophers maintain that phenomenal consciousness and conscious accessibility are distinct phenomena requiring distinct explanations (e.g. Block 2007, 1995). Another influential theory is Tononi's (2004) Integrated Information Theory (developed out of Edelman 2003), which Tononi presents as a theory of phenomenal consciousness, though because it casts consciousness as a kind of causal-informational integration, it rather looks like a theory of access consciousness.

Taking the theories of consciousness that are on the table, the natural approach is to see whether any one of them can also account for conscious unity. This seems possible for any of the major neurofunctional theories, including the two mentioned above. For example, if the entry of some content into the global workspace makes it an element of conscious experience, then perhaps two contents' simultaneous entry into the same global workspace makes their elements unified. Meanwhile, Tononi's (2004) theory of consciousness explicitly appeals to the integration of contents; whatever theory of phenomenal consciousness it offers is quintessentially holistic.

It's striking that these prominent theories of consciousness are so readily amended into theories of conscious unity. This is because they offer accounts of what happens to contents to make them contents of *experience*. They can then attribute experiential "togetherness" to the "togetherness" of experiential contents—to their being made conscious together.

Philosophers have occasionally expressed skepticism toward such content-centric explanations of consciousness. Consider Searle's (2000) contrast between *building block* and *unified field* accounts of consciousness. In the former, a subject's consciousness is constructed from elements understood to be conscious independently of their incorporation into a larger experiential whole. Searle suggests that this sort of account sails past consciousness itself, in pursuit of contingently conscious contents. According to a unified field account, the subject's consciousness itself is prior: specific contents become conscious only by modifying their subject's consciousness. The subject must already *be* conscious in order for their *phenomenal field* to be modified in any way—in order to experience anything at all.

One way of interpreting the unified field account is as taking phenomenal consciousness to be a kind of *creature consciousness* (Rosenthal 1993)—a condition of the entire subject. It cannot be the condition of wakeful alertness however (although this is the way Searle 2000 speaks), since, as Bayne (2007) points out, we are phenomenally conscious while dreaming. We might think of this creature consciousness as a state of readiness to enjoy phenomenal experiences.

It is not clear which questions about phenomenal unity would be resolved by adopting a unified field account. The account rules out the possibility of radical atomism, but does not entail that a subject can only have a single unified field at a time (a point conceded by Bayne 2007). For one thing, to know that what produces an *X* is a mechanism of type *M* isn't yet to know *how many* *Xs* each such mechanism can produce; nor is it to know how many mechanisms of that type there are within the brain.

The deeper issue is that we do think of conscious unity precisely in terms of relations between contents (again, see Bayne 2007). What after all is a *unified field*—the supposedly prior conscious "thing" that is only modified in this way and that by particular conscious contents? Searle essentially defines a unified field as an experiential whole. In that case, stating that two contents modify the same field does not explain their unity but is equivalent to stating that they are unified (as Prinz 2013 notes). If a phenomenal field is instead defined in a way that is neutral with respect to the relations between the contents modifying a field, then it is not clear that a subject with a single phenomenal field will necessarily have a unified consciousness. Why couldn't multiple contents modify their common field *so as to* generate a disunified consciousness?

## 6 The Disunity Question: Can Conscious Unity Fail?

Much of the philosophical literature on the unity of consciousness concerns candidate cases in which it fails. A subject's consciousness could, logically, fail to be unified in any of a number of ways, including the following three. First, a subject with a *radically atomistic* consciousness would have a multitude of elements of experience, none incorporated into any larger whole. In a subject with a *partially unified* consciousness, each element of experience would be unified with some but not all others. This would require that conscious unity not be a transitive relation between experiences or elements of experience. A subject with a *multiple consciousness* would instead possess multiple experiential wholes, not unified with each other, though the elements within each whole would be unified.

Candidate cases of this last sort have received the most philosophical attention. A subject with a multiple consciousness has multiple experiential wholes, rather than a single experiential whole incorporating everything the subject experiences. Most disruptions or impairments of consciousness do not have this characteristic, but instead rather reduce the number of elements encompassed by an experiential whole (simultagnosia, unilateral neglect) or else distort its contents (hallucinations, perhaps Capgras syndrome or apperceptive agnosia). Philosophers have, however, debated several kinds of cases that may be instances of multiple consciousness, most often *dissociative identity disorder* and the so-called *split-brain phenomenon*.

Dissociative identity disorder (DID)—formerly known as “multiple personality disorder”—is one of several dissociative conditions recognized by the Fifth Edition of the Diagnostic and Statistical Manual of Mental Disorders (DSM-5). All such conditions are characterized by impaired integration of consciousness, emotion, and memory, but DID attracts by far the most philosophical attention, because its disruptions and breakdowns are patterned in a unique way. A person with DID might, for instance, appear to suffer a kind of fluctuating amnesia about her gambling habit, but the fluctuations don't appear to be random: she might, say, *always* remember previous episodes of gambling, but *only* whenever she is already gambling or on her way to gamble; at other times she might always be amnesic for such episodes.

Many philosophers believe that this kind of personal and experiential memory is important to remaining the *same person* over time. Fluctuating amnesia might be viewed as compromising the kind of extended diachronic conscious unity that characterizes ordinary human experience. Certainly, one point of contention about dissociative identity disorder (Tye 2003; Braude 1991) concerns whether someone with DID has different, mutually interrupting perspectives or streams of consciousness at different times.

Another reason why many philosophers see DID as a candidate case of multiple consciousness, however, is that the condition is often characterized as one in which multiple persons somehow animate one body (Radden 1996; Humphreys and Dennett 1989). If we understand the condition in this way, it is not that the hypothetical DID subject mentioned above remembers previous episodes of gambling only during occurrent episodes; rather, her body is the body of at least *two* experiencing beings, one of whom gambles however often and naturally remembers having gambled before, and one who has never gambled and therefore has no memories of having done so. If understood in this way, DID appears to be a case of diachronic multiple consciousness within a human being—though not necessarily within an experiencing person (see Section 7 on “identity questions”).

Although dissociative identity disorder is more striking than the split-brain phenomenon, DID has the disadvantage of continuing to be the object of fierce controversy among clinicians and cognitive psychologists (see, for instance, Lynn *et al.* 2014 vs. Dalenberg *et al.* 2014). In contrast, the basics of the split-brain phenomenon are well accepted. The split-brain phenomenon

is a result of medical neurosurgery, colloquially known as “split-brain surgery,” that cuts through the corpus callosum connecting the two cerebral hemispheres. Since the corpus callosum is a conduit for interhemispheric interaction and information exchange, the surgery naturally alters and impedes this interaction and exchange (though in no way prevents it). The split-brain phenomenon is generally discussed as a candidate case of synchronic conscious disunity: are all of a split-brain subject’s experiential elements at every moment incorporated into a single experiential whole, or is each element instead incorporated into *one of two* experiential wholes—one associated with the right hemisphere and one with the left?

The question is difficult to resolve partly for straightforward empirical reasons: the cases are complicated, and while some kinds of experience appear to be cleanly interhemispherically dissociated in some subjects at some times, others do not. A more basic theoretical problem is that any evidence of conscious disunity takes the form of apparent lack of integration amongst experiential elements—yet lack of integration constitutes *prima facie* evidence that the elements were not in fact *conscious* in the first place.

In addition, having a *multiple* consciousness does not seem to be the sort of thing one could introspect. To think otherwise would be to suppose that a subject with a multiple consciousness could compare multiple elements of experience and judge that they were not elements of one whole. But then how would she have made the comparison? Any introspective act will itself be an element of some experiential whole or other, and can survey only the elements of that whole.

If introspection cannot reveal failures of phenomenal, access, or awareness unity, then even a neurotypical human being cannot know *via* introspection that her consciousness is perfectly unified, or even unitary. (See Marcel 1993 on the potential disunity of neurotypical consciousness over very short time scales; see also Dennett 1991.) Among other things, we could be subject to a version of the so-called *refrigerator light illusion*, if something about the very act of introspecting the elements of our experience itself unifies them (Prinz 2013). Perhaps consciousness is disunified whenever we aren’t looking!

## 7 Identity Questions: Selves, Self-Consciousness, and Subjective Perspectives

From one perspective, it’s hardly radical to suppose that neurotypical consciousness may not be wholly unified or even unitary. We already know that our minds are not wholly rationally unified, self-knowing, or self-controlled, and that introspection can mislead about all manner of things. Scientists and philosophers recognize that human reasoning and perception are subserved by a multitude of systems, perhaps even mind-like systems (Evans and Frankish 2009). So why not division or multiplicity within consciousness itself? If we feel that we have a better claim to our *phenomenal* unity, perhaps this is just pride speaking; perhaps it is just because our phenomenal properties are dimly felt to lie within the last realm that science has yet to conquer (Dennett 1988: 386).

On the other hand, phenomenal unity may be unique in the closeness of its connection to our first-personal ways of thinking about the identities of experiencing subjects. This connection is mediated by the concept of a subjective perspective, and it poses obstacles to understanding phenomenal disunity that we don’t face when thinking about psychic disunity of other kinds.

Experiencing beings have perspectives: there are facts about the way the world *is*, and further facts about the way the world *appears to them*. Rocks do not have perspectives. There are no facts about the way the world appears to be to a rock; there is just the way the world is, including that portion of the world that is the rock.

Our first-personal way of understanding the conditions of being an experiencing subject is in terms of having a perspective. It is first-personal because the concept of a subjective perspective is itself one that we grasp only because each of us already has a perspective (Nagel 1974). We might conceptualize what it is for a being, X, to have a subjective perspective, in terms of there being something that would count as successfully imagining being X. For an agent Y at time  $t_y$  to successfully imagine being X at time  $t_x$  is for Y to (intentionally) experience at  $t_y$  all and only what X experiences at  $t_x$ . These success conditions are very stringent—presumably impossible to meet—but this is consistent with X's having a perspective. What is required for X to have a perspective is that an attempt to imagine being X has success conditions. In contrast, nothing would count as successfully imagining being a rock.

Now suppose that at some time  $t_p$  some target being, T, does not have a single experiential whole, but rather has two of them, W1 and W2: the elements of W1 are all unified with each other, but not with those of W2, and analogously for the elements of W2. And suppose that agent Y wants to imagine at time  $t_y$  being T at  $t_p$ . Suppose that Y is fantastically good at such acts of empathic imagination: given enough third-personal information about a target, Y can undergo precisely those experiential elements undergone by Y's target. On the other hand, suppose that Y has perfect control over only the *contents*, and not the *structure*, of Y's own experience. So whatever the elements of Y's experience, these elements are always synchronically phenomenally unified, whether Y is trying to imagine being someone else or not.

Y then cannot successfully imagine being T-at- $t_p$  by engaging in only a single act of empathic imagination. If Y chooses to undergo at  $t_y$  every experiential element of T at  $t_p$ , these elements would all be phenomenally unified, since this is the fixed structure of Y's consciousness. Y would thus experience everything T experienced at  $t_p$ , but not *only* what T experienced, since T does not experience unity between the elements of W1 and those of W2, and Y would. If Y instead chooses to undergo at  $t_y$  only every element of, say, W1 at  $t_p$ , and no element of W2, Y would thus experience *only* what T experienced at  $t_p$ , but not *everything* T experienced.

One could debate whether Y could succeed at occupying T-at- $t_p$  via two successive imaginative acts, in which Y undergoes at  $t_y$  the elements of W1 and undergoes at  $t_{y+1}$  the elements of W2 (see Bayne 2010 and Schechter 2018). But it is at least clear that there is no other way for Y to succeed.

This is because T does not have *a* subjective perspective, singular, for Y to imaginatively take on. Rather T has *two* perspectives: W1 and W2. The success conditions for Y's imagining being T must be relativized to W1 and W2: there are success conditions for imagining being T-subject-to-W1 and T-subject-to-W2. But recall that these are conditions on imagining *being someone*. So now it begins to seem to us—and should to Y—that T is somehow *two* subjects of experience.

The connections between phenomenal unity and subjects of experience bear on the assumption that phenomenal unity is a transitive relation. Suppose that it weren't, and that two experiences not unified with each other, E1 and E2, were both unified with a third, E3. Because they are not unified with each other, E1 and E2 should be elements of distinct perspectives. We have seen that where there are distinct perspectives, P1 and P2, there is pressure to posit distinct subjects of them, S1 and S2. Perspectives presumably incorporate every element of experience that is unified with any element they do incorporate. Since E3 is unified with P1's E1 and with P2's E2, E3 should be an element of both P1 and P2. S1 and S2 should thus both be subjects of E3—yet if elements of experience get their identities partly from the subjects whose experiences they are, a single element of experience cannot belong to multiple subjects.

This contradiction creates pressure to insist that phenomenal unity must after all be transitive. Indeed the transitivity of (synchronic) phenomenal unity has something close to the status of an axiom in the literature. Lockwood's work (1989) defending the possibility of non-transitive

phenomenal unity is perhaps both the major exception to this rule and the one that proves it, since in later writing he admitted to being unsure himself whether the idea of non-transitive phenomenal unity really makes sense (Lockwood 1994).

Because the identity conditions for subjects of experience and for subjective perspectives are so closely related, there will always be compelling reasons to deny that a human being has a disunified and especially a multiple consciousness (Bayne 2010). Again, recognizing multiple subjective perspectives within a single human being creates pressure to posit distinct subjects of each perspective. But subjects of experience are objects of moral concern, and it is not clear what it would mean for there to be multiple distinct objects of moral concern within a single living being. There will thus always be reason to deny that a single human being is more than a single subject of experience—whatsoever the empirical facts about consciousness may turn out to be.

## References

- Baars, B. (1988) *A Cognitive Theory of Consciousness*, Cambridge, UK: Cambridge University Press.
- Bartels, A. and Zeki, S. (1998) “The theory of multistage integration in the visual brain,” *Proceedings of the Royal Society of London B (Biological Sciences)* 265: 2327–2332.
- Bayne, T. (2007) “Conscious states and conscious creatures: Explanation in the scientific study of consciousness,” *Philosophical Perspectives* 21 (*Philosophy of Mind*): 1–22.
- Bayne, T. (2010) *The Unity of Consciousness*, Oxford: Oxford University Press.
- Bayne, T. and Chalmers, D. (2003) “What is the unity of consciousness?” In A. Cleeremans (ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*, Oxford: Oxford University Press.
- Block, N. (1978) “Troubles with functionalism,” in C. Savage (ed.), *Perception and Cognition: Issues in the Foundations of Psychology*, Minneapolis, MN: University of Minnesota Press.
- Block, N. (1995) “On a confusion about a function of consciousness,” *Behavioral and Brain Sciences* 18: 227–287.
- Block, N. (2007) “Consciousness, accessibility, and the mesh between psychology and neuroscience,” *Behavioral and Brain Sciences* 30: 481–548.
- Braude, S. (1991) *First Person Plural: Multiple Personality and the Philosophy of Mind*, New York: Routledge.
- Bennett, D. and Hill, C. (2014) “A unity pluralist account of the unity of experience,” in D. Bennett and C. Hill (eds.), *Sensory Integration and the Unity of Consciousness*, Cambridge, MA: MIT Press.
- Brook, A. and Raymond, P. (2014) “The Unity of Consciousness,” in E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/win2014/entries/consciousness-unity/>
- Carruthers, P. (2017) “Block’s overflow argument,” *Pacific Philosophical Quarterly* 98: 65–70.
- Chisholm, R. (1981) *The First Person: An Essay on Reference and Intentionality*, Minneapolis, MN: University of Minnesota Press.
- Church, J. (1995) “Fallacies or analyses?” *Behavioral and Brain Sciences* 18: 251–252.
- Dainton, B. (2000) *Stream of Consciousness*, London: Routledge.
- Dainton, B. and Bayne, T. (2005) “Consciousness as a guide to personal persistence,” *Australasian Journal of Philosophy* 83: 549–571.
- Dalenberg, C., Brand, B., Lowenstein, R., Gleaves, D., Dorahy, M., Cardeña, E., et al. (2014) “Reality versus fantasy: Reply to Lynn et al. (2014),” *Psychological Bulletin* 140: 911–920.
- Dehaene, S. and Naccache, L. (2001) “Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework,” *Cognition* 79: 1–37.
- Dennett, D. (1988) “Quining qualia,” in A. Marcel and E. Bisiach (eds.), *Consciousness in Contemporary Science*, Oxford: Clarendon Press.
- Dennett, D. (1991) *Consciousness Explained*, New York: Little, Brown.
- Dennett, D. (2001) “Are we explaining consciousness yet?,” *Cognition* 79: 221–237.
- Edelman, G. (2003) “Naturalizing consciousness: A theoretical framework,” *PNAS* 100: 5520–5524.
- Evans, J. and Frankish, K. (2009) “The duality of mind: A historical perspective,” in J. Evans and K. Frankish (eds.), *In Two Minds: Dual Processes and Beyond*, Oxford: Oxford University Press.
- Fodor, J. (1981) “The mind-body problem,” *Scientific American* 244: 114–125.
- James, W. (1890) *The Principles of Psychology*, New York: Dover.

- Humphreys, N. and Dennett, D. (1989) "Speaking for our selves: An assessment of multiple personality disorder," *Raritan* 9: 68–98.
- Hurley, S. (1998) *Consciousness in Action*, Cambridge, MA: Harvard University Press.
- Levine, J. (1983) "Materialism and qualia: The explanatory gap," *Pacific Philosophical Quarterly* 64: 354–361.
- Lockwood, M. (1989) *Mind, Brain and the Quantum*, Oxford: Blackwell Publishers.
- Lockwood, M. (1994) "Issues of unity and objectivity," in C. Peacocke (ed.), *Objectivity, Simulation, and the Unity of Consciousness*, Oxford: Oxford University Press.
- Logothetis, N. (1998) "Single units and conscious vision," *Philosophical Transactions of the Royal Society of London Series B* 353: 1801–1818.
- Lynn, S. J., Lilienfeld, S. O., Merckelbach, H., Giesbrecht, T., McNally, R., Loftus, E., et al. (2014) "The trauma model of dissociation: Inconvenient truths and stubborn fictions. Comment on Dalenberg et al. (2012)," *Psychological Bulletin* 140: 896–910.
- Marcel, A. (1993) "Slippage in the unity of consciousness," in G. Bock and J. Marsh (eds.), *Experimental and Theoretical Studies of Consciousness*, Chichester, UK: John Wiley & Sons.
- Marks, C. (1981) *Commissurotomy, Consciousness and Unity of Mind*, Cambridge, MA: MIT Press.
- McGurk, H. and MacDonald, J. (1976) "Hearing lips and seeing voices," *Nature* 264: 746–748.
- Nagel, T. (1971) "Brain bisection and the unity of consciousness," *Synthese* 22: 396–413.
- Nagel, T. (1974) "What is it like to be a bat?," *Philosophical Review* 83: 435–450.
- Prinz, J. (2013) "Attention, atomism, and the disunity of consciousness," *Philosophy and Phenomenological Research* 86: 215–222.
- Radden, J. (1996) *Divided Minds and Successive Selves: Ethical Issues in Disorders of Identity and Personality*, Cambridge, MA: MIT Press.
- Revonsuo, A. (1999) "Binding and the phenomenal unity of consciousness," *Consciousness and Cognition* 8: 173–185.
- Rosenthal, D. (1986) "Two concepts of consciousness," *Philosophical Studies* 49: 329–359.
- Rosenthal, D. (1993) "State consciousness and transitive consciousness," *Consciousness and Cognition* 2: 355–363.
- Schechter, E. (2013) "The unity of consciousness: Subjects and objectivity," *Philosophical Studies* 165: 671–692.
- Schechter, E. (2018) *Self-Consciousness and "Split" Brains: The Minds' I*, Oxford: Oxford University Press.
- Searle, J. (2000) "Consciousness," *Annual Review of Neuroscience* 23: 557–578.
- Shoemaker, S., (1996) "Unity of consciousness and consciousness of unity," in *The First-Person Perspective and Other Essays*, Cambridge: Cambridge University Press.
- Tononi, G. (2004) "An information integration theory of consciousness," *BMC Neuroscience* 5: 42.
- Treisman, A. (1998) "Feature binding, attention, and object perception," *Philosophical Transactions: Biological Sciences* 353: 1295–1306.
- Treisman, A. and Gelade, G. (1980) "A feature integration theory of attention," *Cognitive Psychology* 12: 97–136.
- Tye, M. (2003) *Consciousness and Persons: Unity and Identity*, Cambridge, MA: MIT Press.
- Zihl, J., Von Cramon, D. and Mai, N. (1983) "Selective disturbance of movement vision after bilateral brain damage," *Brain* 106: 313–340.

## Related Topics

Consciousness, Personal Identity, and Immortality  
 Multisensory Consciousness and Synesthesia  
 Consciousness and Psychopathology  
 Consciousness, Free Will, and Moral Responsibility  
 Representational Theories of Consciousness  
 The Global Workspace Theory  
 The Multiple Drafts Model  
 Consciousness and Attention  
 Consciousness, Time, and Memory