



Phenomenology and the unity of consciousness

Graham Peebles¹

Received: 11 October 2019 / Accepted: 13 January 2021 / Published online: 27 January 2021
© Springer Nature B.V. 2021

Abstract

The phenomenology of the unity of consciousness can be analyzed in terms of perceptual spatial and object unity. Subject unity—what we commonly understand by “the unity of consciousness”—has no attendant phenomenology. The further, non-phenomenological, effects of unity can be analyzed in terms of the functional notion of access unity. The unity of consciousness in general can therefore be analyzed in terms of access unity. As a consequence, we can avoid the theoretical introduction of problematic notions such as subsumptive or phenomenal unity.

Keyword Unity of consciousness · Consciousness · Functionalism

1 Introduction

The longer one considers the phenomenology of the unity of consciousness the more puzzling it can become, for although it may be more or less clear what is involved in the phenomenology of, for example, seeing a red square, the more one ponders the phenomenology of seeing a red square and a yellow circle, or seeing a red square and a yellow circle and hearing a low hum, or seeing a red square and a yellow circle, hearing a low hum, and feeling a pain in one’s finger while worrying about it, all together as a “unified experience” or as “unified experiences”, the less sure one’s introspective footing can become. Some will contend, with intuitive merit, that there is a phenomenology of unity which resists a factoring into the individual states. Others will contend that such a factoring is possible without phenomenological remainder. My sympathies lie in the direction of the latter view, and I would like to make

✉ Graham Peebles
peebles@ams.eng.osaka-u.ac.jp

¹ Intelligent Robotics Laboratory, Graduate School of Engineering Science, Osaka University, 1-3 Machikaneyama, Toyonaka, Osaka 560-8531, Japan

the case in this paper for the plausibility of this relatively unpopular view. This view of unity is minimal and its phenomenology is sparse; indeed, one could describe it as “eliminativist” about the phenomenology of unity.¹

My discussion will be conducted within the framework of a reductive first-order functional-representational theory of consciousness,² as the theoretical resources employed are only those of representational content and functional access, and a goal of this paper is to bolster the plausibility of this theory of consciousness.³ On a first-order functional-representational theory of consciousness, if there is a phenomenology to unity which outstrips the phenomenology of the unified experiences, each of which can be individually analyzed representationally, this can seem to call the general theory into question. The theory has difficulty in dealing with such alleged aspects of phenomenology as the unity of consciousness, because it is difficult to analyze this phenomenology in terms of first-order contents.

However, I would like to argue that there is, in fact, no phenomenology to the overall unity of consciousness, or subject unity. The phenomenology of the unity of consciousness can be accounted for solely within perceptual unity, where two objects are perceived as being in the same physical space or two properties are perceived as being properties of the same object. The “top level”, so to speak, of unity—what we commonly understand by “the unity of consciousness”—has no attendant phenomenology which outstrips that of this perceptual spatial and object unity. The behavioural and cognitive effects of subject unity can be analyzed in terms of the functional notion of access unity without the need for further phenomenology. There is no introspective or theoretical need to introduce a phenomenological notion of unity such as those defended in different ways by, for example, Bayne (2010), Bayne and Chalmers (2003), Dainton (2000), Shoemaker (2003) and Tye (2003). My case is, it could be put, a positive one as opposed to a negative one. I will argue for the view by motivating it and then defending it from the major objections. This speaks against opposing views indirectly, by making the case that one need not accept the series of arguments and intuitions that can lead in that direction.

In Sects. 2 and 3, I will discuss subject unity and the phenomenology of object and spatial unity. In Sect. 4, I will defend the claim that access unity can account for the unity of consciousness from the major argument against it that originates in

¹ The arguments in this paper are confined to synchronic unity. Diachronic unity is a further question.

² I mean here theories which understand consciousness to be, in the first place, a matter of the processing of representations with first-order (i.e. external, non-mental) content. Examples of this type of theory include Baars (1989), Dennett (1991), Dretske (1995), Prinz (2012). The motivating arguments do not depend on an acceptance of this theory. One could agree that there is no phenomenology to unity per se, yet reject the further claim that the unity of consciousness can be analyzed in terms of access unity. Hill (1991, 2014, 2018), for example, rejects a notion of a phenomenology of unity but does not analyze unity in general as access unity. Furthermore, I would reject more speculative types of phenomenology, such as cognitive phenomenology, or representational theories such as Shoemaker (1994) on which theoretical room is made for properties of experience, which could, perhaps, be employed to deal with a phenomenology of the unity of consciousness, although the theory presented is not incompatible with these theories.

³ Without recourse to controversial notions of unity such as that proposed in Tye (2003). See Sect. 5 for discussion.

Bayne and Chalmers (2003). In Sect. 5, I will defend the view from the objection from the breakdown of perceptual spatial unity. Finally, in Sect. 6, I will discuss some related aspect of other accounts, extend my argument to the general notion of a conscious field, and briefly discuss some further questions about the unity of consciousness.

2 Subject unity

It is intuitively easy to get a handle on what is meant by saying that a subject's consciousness is unified, or that their conscious states are unified. Here, for example, is Bayne describing an episode in his stream of consciousness at the outset of an influential recent book on the subject.

I'm sitting in the Cafe Cubana (47 Rue Vavin, Paris). I have auditory experiences of various kinds: I can hear the bartender making a mojito; I can hear the dog behind me chasing his tail; and there's a rumba song playing somewhere on a stereo. I am enjoying visual experiences of various kinds: I can see these words as they appear in my notebook; I can see the notebook itself; and I have a blurry visual impression of those parts of the room that lie behind the notebook. Co-mingled with these auditory and visual experiences are olfactory experiences of various kinds (I can smell something roasting in the kitchen); bodily sensations of various kinds (I am aware of my legs under my chair; I can feel my fingers on the table); and a range of cognitive and affective experiences. The bartender is talking to an old woman at the bar, and I have a vague sense of understanding what he's saying. I am soon to embark on a lengthy trip, and a sense of anticipation colours my current experiential state. Finally, I am enjoying conscious thoughts. I realize that the bar is about to close, and that I will be asked to leave if I stay for much longer (2010, p. 5).

We are all familiar with such an episode. Indeed, our stream of consciousness is at least in significant part composed from such episodes. Later, Bayne picks out the key element.

Consider again what it's like to hear a rumba playing on the stereo whilst seeing a bartender mix a mojito ... There is something it is like to hear the rumba, there is something it is like to see the bartender work, and there is something it is like to hear the rumba while seeing the bartender work. Any description of one's overall state of consciousness that omitted the fact that these experiences are had together as components, parts, or elements of a single conscious state would be incomplete. Let us call this kind of unity—sometimes dubbed 'co-consciousness'—phenomenal unity (2010, p. 10).

The notion of co-consciousness or phenomenal unity is a familiar one from the literature (Lockwood 1989; Hurley 1998; Dainton 2000; Shoemaker 2003; Tye 2003), and although it is, of course, a significant task to provide an analysis of co-consciousness, it is intuitively easy to grasp the basic notion. My perceptions of the colour of my laptop and the shape of my laptop are co-conscious, or unified. These perceptions are also unified with my wondering where to eat tonight. It is perhaps less intuitive that this wondering is unified with my perception of my laptop than that my perceptions of the different properties of the laptop are unified with each

other, but there is a clear contrast which brings out that they are: the unity of my perception of my laptop and thought about where to eat tonight, and the disunity of my perception of my laptop and your thought about where to eat tonight. The latter are disunified whereas the former are unified.

Such observations do not get us very far in answering any significant questions about the unity of consciousness, but they do point to an idea that one is naturally drawn to: that there cannot be unity of consciousness between subjects. That is, two conscious states cannot be unified if they are states of different subjects. This seems to provide us with a conceptual link between the notions of unity and of a subject. Perhaps, then, given that it is tempting to say that it is impossible for there to be unity across subjects, it is also impossible for there to be disunity within a subject. If there were to be such disunity, then there would be two subjects. Exactly what to say about the modal strength of unity or the relation between subjects and the unity of consciousness is a problem of the utmost difficulty which likely turns in the end on what we say about the oddest of empirical cases (see Bayne 2010, Part II for extended discussion). These cases are the most controversial and difficult, and so thankfully I do not need to take a stand on the questions involved therein. I will make no claim about the modal nature of the unity of consciousness, nor will I discuss whether partial unity is possible. My focus is on the functional and phenomenological consequences of unity and the theoretical resources required in its analysis.

Subject unity is, one might think, a trivial notion. To say that two states are subject unified is not to say much at all. It is, though, to say something—after all, my thoughts are not subject unified with yours—but we seem to intuitively want to say more. There are two aspects that one looks to further explain. One is the cognitive and behavioural effects of unity. The other—what seems to occupy most philosophers—is an account of the phenomenology of subject unity. Typically, this motivates a special, phenomenological type of co-consciousness, or phenomenal unity, in terms of which subject unity is analysed. The conclusions about the nature of the unity of consciousnesses that philosophers have drawn from this starting point differ, but they mostly share the intuitive starting point that there is, to speak somewhat metaphysically loosely, an entity that is the unity that there is a phenomenology to such unity. Bayne's influential account, for example, is based on the notion of 'subsumptive unity' in which a complex state subsumes the unified states as in a part-whole relation. Tye's (2003) analysis is in terms of a single conscious state. Dainton's (2000) analysis is in terms of a relation of co-consciousness between the unified states.

These theories differ in their metaphysics of unity, but they share the aforementioned metaphysical and phenomenological starting point, namely that there is a phenomenology of unity per se and that individual unified states bear a relation to the unity that is akin to the part-whole relation. I would like to argue that the phenomenology of unity can be accounted for in terms of the content that results from spatial and object unity—the unities that lead to experiences of properties being unified as experience of objects and of these experiences being unified as experiences of objects in the same space. The further aspects of unity—the non-phenomenological aspects of spatial and object unity, and the non-phenomenological unities of other states—can be analyzed in terms of access unity. Access unity is the unified

accessibility of those states for the rational control of speech, thought, and action. Unity in general is access unity and therefore unity is both non-phenomenological and non-introductory of a total experience, conjoint experience, phenomenal field, or subsumptive experience, etc.; non-introductory of any entity over and above the functional network of unified experiences. Where one does find a phenomenology of unity within the network of access-unified experiences, this can be explained by the perceptual object and spatial content that results from the object and spatial unity of the experiences in question.

An initial motivation for this view is introspective: when I introspect, I do not find a phenomenology of unity *per se*. Rather, I only find phenomenology of unified perceptual experiences where this results in the contents of spatial and object unity. This allows for unity to be analyzed in terms of a relation which is non-phenomenological in itself, suitably unifying, and which can introduce the contents of perceptual spatial and object unity. Access-unity is such a relation. This can be seen as support for a first-order functional-representational theory, as it is access unity in terms of which this theory of consciousness trades.

One can find views with similar aspects in the literature. Hill (1991, 2014, 2018; Bennett and Hill 2014), for example, argues for a “multiple relations” account of unity on which there are a variety of relations which unify experiences, but not a general unity relation which unifies all unified experiences and is phenomenological or ontologically introductory in that it introduces an overall or joint experience over and above the related unified experiences themselves. Hill argues that introspection does not reveal a phenomenology to unity and that unity can be theoretically analyzed in terms of relations between the unified states without the introduction of the experiential analog of the whole to its parts. The view that I present differs significantly from Hill’s in that, on this account, unity is analyzed in terms of access unity. There are similarities, though, in that access unity is neither phenomenological in itself even though in perceptual cases of object and spatial unity it results in the introduction of content and thus perceptual phenomenology, and nor is it ontologically introductory.⁴ Masrour (2014) likewise rejects such unity relations, and argues for a “connectivity view” on which “we can account for phenomenal unity in terms of experiences of specific relations” (2014: 328). On this view, there are relations between the individual experiences in terms of which they are unified, rather than the experiences being members of a mereological whole, single experience, etc. While I agree with this, Masrour’s theory differs significantly in that he posits both a variety of unifying relations between experiences and also that it is our experience of these relations which unifies the experiences. On the view presented here, the unity relation is access unity and therefore is not phenomenologically manifest to introspection. Only the contents that are introduced in spatial and object unity are so manifest.⁵

⁴ It is part of Hill’s theory that unity can be partial and that there can be disunity. As mentioned above, my account is neutral on these questions.

⁵ Let me emphasize a clarification. I am not arguing that all unity is spatial or objectual. This is more or less an indefensible position as there are cases of unity in which there is no plausible spatial or objectual

This is not necessarily to say that the rejected notions in other theories of unity ought to be considered contradictory, incoherent, or lead to unacceptable consequences. Rather, I would like to press the case that we need not introduce them and then defend an account of unity in terms of access-consciousness on which they are not required, draw out its consequences, and defend it from objections. My argument is therefore not one of elimination, but, rather, it is a proposal that the view presented is worthy of consideration. Let me now move on to discuss spatial and object unity, and then access unity, in some detail.⁶

3 Spatial and object unity

Consider a perception of a red square and a green circle. We have two experiences of properties bound into different discrete objects: experiences of the red colour and of the square shape as the properties of one object and of the green colour and the circular shape as the properties of another. The respective pairs of properties are unified with each other in that they have object unity, or are object unified. Were they to lack this unity, the overall experience would be of something like disconnected, disparate property overlays, as opposed to the discrete objects with which we are familiar. It is therefore clear that perceptual object unity does result in some phenomenology. This is because it results in some difference in the contents of perception. Consider how perceptions of property F at location l and of G also at l would seem when object disunified as opposed to unified: that there is an F object at l and a G object at l ; that there is an F object at l , a G object at l , and the two objects are one and the same—i.e. that there is an $F G$ object at l or that the object at l is an FG . It does not matter how we put this in natural language for the present point, only that we recognize that object unity results in a difference in content of this kind.⁷ The phenomenological effect of object unity can be understood in terms of

Footnote 5 (continued)

content (see, e.g., Roelefs 2014 for a list). Rather, I am arguing that the phenomenology of unity can be accounted for in terms of the contents of object and spatial unity within perception.

⁶ One could ask whether the account is eliminativist, a term I scare-quoted above, or reductionist about the unity of consciousness and the phenomenology of unity. If one can helpfully distinguish between these two notions for all individual cases, the most natural reading is that the account reduces unity to access unity, reduces the phenomenology that does result from unity to the contents of perceptual spatial and object unity that it introduces, and eliminates a phenomenology of unity per se.

⁷ As a result of properties being bound together in sub-personal processing, they are consciously represented differently than they would be were they delivered to conscious experience unbound. In a discussion of the binding problem, Garson (2001) argues persuasively that once the binding problem is solved at the physical level of implementation, there is no further phenomenological problem that remains. I find such an argument quite convincing. Binding has a representational consequence—this is what binding does—and that is all that there is to say. Why would there need to be phenomenological binding over and above representational binding? The same point applies to “functional binding”. Once two functional properties are appropriately located within the same functional nexus, what more would be required for unity? There remains the tricky question just alluded to as to exactly how bound contents articulate the representation of objecthood. However, this is not really a problem of “phenomenal binding”, as each option equally well binds, but of carefully describing our conceptual scheme as it appears to us in experience. For example, the two formulations used above—that there is an $F G$ object at l or that the object at

this resulting content. Within perception, this phenomenological effect also obtains for spatial unity, where two states are spatially unified just when perception presents their contents as being located in the same space.

The phenomenology of spatial and object unity, it seems to me, is to be found in the network of unified perceptual states. These cases of unity result in the contents of the unified experiences being unified in such a way as to introduce some further perceptual content. This reflects the results of sub-personal spatial and object binding. Within perception, therefore, these unified experience result in a phenomenology that goes beyond the unified contents considered individually. One experiences the shape of an object and the colour of an object; however, if these experiences are object unified, then one also experiences a discrete object with both this shape and this colour. Likewise for spatial unity. For example, I experience the colour of my laptop and also its shape. However, as these come to consciousness not just unified but object unified, I also experience my laptop as a discrete object with this colour and shape. The same point holds for my desk, and holds with respect to the spatial relation between the two. It is within this network of perceptual object and spatial unities, I contend, that all of the phenomenology of unity is to be found. No other cases of unity result in such a phenomenology because they do not result in some further content, such as that which presents spatially related discrete objects in perception. What holds unified experiences together in general, as I will discuss in more detail in Sect. 4, is access unity. And, unlike perceptual content that results from binding in the aforementioned cases of unity, access unity itself is not phenomenological.

Object and spatial unity extend to cognition. However, here intuition and introspection seem to diverge. When I think about the table in the corner of my office, for example, I think of two discrete objects (allow me to include the office as an “object”) standing in a spatial relation. Unlike the intra-perceptual case, however, I do not believe that it is proper to talk of there being phenomenology to this unity. Likewise for such cases as thinking about putting a lamp into the corner of my office while looking at the corner to consider whether or not it is a good idea. The two states are spatially unified. However, the unity of these two states does not seem to me to have any phenomenology. That is, there is no phenomenology *qua* the spatial unity of these two states.

The majority of philosophers, though, seem to reject this claim as counterintuitive. One could dispute an aspect of this by, for example, holding that there is cognitive phenomenology, i.e. that thinking that p has phenomenology of the same type as perceiving that p and that there is a phenomenology of unity across the perception-cognition boundary. This is not the place to discuss this in detail, and so I will simply state that I reject the notion of cognitive phenomenology in addition to these

Footnote 7 (continued)

I is an F G —are different in their logical construction and therefore in their representation of the basics of our conceptual scheme as regards the structure of the external world. We must, therefore, be careful when it matters which formulation is the correct one, such as when we are considering questions about the philosophy of language or mental content, or some issues in metaphysics. But these questions are not relevant here.

alleged cases of a cross-boundary phenomenology of unity.⁸ A further option that one might consider is that there is a unifying cognitive awareness of all conscious states as one finds suggested, for example, in higher-order thought theories of consciousness such as Rosenthal (2005) or reflexive theories such as Kriegel (2009). On such theories, the subject is, for example, (unconsciously) aware of all of their experiences as theirs (higher-order) or reflexively consciously aware of their own experiences (reflexive) and one can imagine a theory of unity which makes use of these resources to analyze unity. This is not the place to discuss this in detail, but I would reject these theories of consciousness and both are also compatible with the presented theory as both theories can be framed entirely in terms of access consciousness.

Major proponents in the unity literature tend, though, not to extend the type of unity that one finds between the contents of perceptual experiences to other conscious states, nor to advert to a higher-order state of global awareness. Rather, they tend to introduce a *sui generis* phenomenological relation that holds between unified conscious states (Lockwood, Dainton, Bayne, being major examples). This is my focus. Set aside, however, the specific analyses of the phenomenon and simply consider again whether there really is such a phenomenon at all that requires analysis. It seems to me not.

When I perceive the corner of my room and think about putting a lamp there, there just does not seem to me to be a *phenomenology to the unity* of the thought and the perception. Likewise, when I engage in perceptual memory or imagination, states in which there is indeed phenomenology as the result of spatial and object unity, when these are unified with perceptual states or with each other there is no phenomenology to this unity. There is more going on, certainly, because there are more states, or collections of states, but there does not seem to be something phenomenological extra that results from these additions. The phenomenology of unity is contained within a phenomenologically encapsulated collections of perceptual states.⁹

It must be noted that the discussion turns significantly on the introspective or intuitive assessment one gives of the claim that the phenomenology of unity is so

⁸ It is also not clear if and how the acceptance of cognitive phenomenology would bear on the present discussion, as the relevant spatial and object unities seem to me to occur within perception and thought, not across the boundary. I, at least, cannot come up with a case where a perception and a thought could be said to be phenomenologically unified in the paradigm intra-perceptual way. One may argue that cases of induced or acquired seeing-as might do the trick, but these can be handled as follows: a conscious thought may cause an experience to acquire extra content, but this extra content does not link the thought and experience in the way that spatial and object unity links the contents within perception through further perceptual content. In addition, cross or multi-modal experiences would also seem likewise to raise no such problem. However, if one insisted on there being a phenomenology of thought, then one could hold that the extra content that results in cognitive analogs of perceptual binding are cases of a phenomenology of unity within thought. Nevertheless, I do not think it proper to talk of a phenomenology of thought, although I will not press the point here as this does not affect my main point that there is no phenomenology to unity per se, only to the content introductions within perceptual unity.

⁹ I will refer to object and spatial unity as “perceptual”, keeping the qualification that this also covers phenomenologically encapsulated perceptual imagination and memory implicit.

encapsulated. This is not a conceptual or metaphysical claim, or even an empirical scientific claim, but a mere psychological claim. When I consider the unity of my perception of the laptop in front of me and my recollection of the restaurant I ate in last night, I simply do not experience any phenomenology *qua* their unity. Matters are clearly different in the intra-perceptual case where, for example, my perceptions of my laptop and desk are spatially unified. This comparison is a distant one, but it serves well to bring out the two types of case. On the one hand we have clear cases of phenomenology of unity within perception, and on the other clear (to my mind) cases where there is no phenomenology of unity. Each case that I have considered either by introspecting my own experience or as a hypothetical example seems to me to fall into one group or the other.¹⁰ If the reader disagrees with this claim, and holds that there is a phenomenology of unity in cases other than intra-perceptual spatial and object unity, then there is not much that I can say in opposition. However, the converse is also true, and so we are at something of an impasse. This impasse is brought about, in part, I believe, by the notorious vagueness inherent in discussing consciousness. Talk of “what it is like”, for example, is perhaps useful as a preliminary, but soon reveals its theoretical imprecision and inadequacy in cases such as the present one. While there is, on a loose interpretation, “something it is like” to perceive *x* and think that *p* in a unified way, this does not really tell us anything substantive about our overall conscious state in this case. Examples such as those given above are, it seems to me, more useful. And when I, at least, compare the intra-perceptual case and the perception-cognition case, I find that there is a phenomenology of unity, manifest in the contents of spatial and object unity, in the former but not the latter case.

Despite the intuitive impasse that arises here, though, there is a great deal more that can be said. Indeed, it is more fruitful to consider the consequences of each side, rather than trade intuitions. For example, one could directly attack the consequences of introducing a *sui generis* phenomenological connection between states, or one could argue that the rejection of such a relation between unified states leads to theoretical difficulties. My concern here is not with the former tactic, as my main aim is not necessarily to refute the opposing view, but with the latter. My main aim is to motivate and defend the sparse view of the phenomenology of the unity of consciousness and then analyze unity in terms of access consciousness. I will now consider the main objection to this view, originating in Bayne and Chalmers (2003).

¹⁰ One might argue that introspection itself could provide a counterexample to this claim. When one introspects, one knows that any two introspected experiences are one’s experiences, but this knowing is not based on a recognition of a phenomenological feature of one’s stream of consciousness but rather a simultaneous noticing of two conscious states. And this noticing is, in turn, just another conscious state that one is in—one whose unity with the others does not result in any phenomenology.

4 Access unity

If we consider the network of unified conscious states within a subject at one time, we can think of a particular group of these states as representing a network of physical objects, including the body and its parts, standing in various spatial relations to one another.¹¹ The perceptual subgroup of these states forms a network of experiences of spatial and object unities whose contents accounts for the phenomenological result of the unity of consciousness. The further effects of unity can be analyzed in terms of access unity. In Block's (1995) famous terms, access consciousness is the immediate availability of the content of a state for the rational control of thought, speech, and action and the account of unity that I would like to defend is that the unity of consciousness is such a functional unity.

There is, though, a significant argument against the suitability of access unity for this role. Bayne and Chalmers (2003) argue that access unity requires that if experiences of p and of q are unified, then the conjunctive content p and q must be available for the immediate rational control of thought, speech, and action. They employ the Sperling experiment as an illustrative example of the pervasive phenomenon of the access bottleneck which is taken to show that p and q is not always so accessible in clear cases of unity between experiences of p and of q . In the experiment, the subjects are briefly flashed a grid of three rows of four letters and afterwards a tone sounds which prompts the subjects to report on the contents of one or more rows. When prompted to report on a single row, the subjects in the original experiment averaged 3.3 correct letters per row of four letters. However, when prompted to report on the contents of the entire grid, their average success rate dropped to 4.5 correct letters out of the grid of twelve. The conclusion is that while the subjects were (highly) access conscious of p and of q (the contents of each of the rows individually) they were not (as) access conscious of p and q (the conjunctive content of the rows). This argument is premised on the acceptance of something like subsumptive unity, although it is perhaps not quite subsumption as there need not be an overall state with the content p and q which subsumes the individual states. And we should not view the accessibility of p and q as the accessibility of the content of a further state over above those with the contents p and q . There is, though, something over and above the accessibility of p and of q on this account, namely the accessibility of p and q . As Bayne and Chalmers note, one could reject this commitment to the accessibility of p and q as a condition on the access unity of the two states—but only, they say, at the cost of 'denying the strong intuition' that there is phenomenal unity and thereby endangering the substantivity of the account. However, I believe that the failure of the accessibility of p and q in these cases does not show that there is no access unity of the experience of p and of q , nor does it endanger the substantivity of the account.

A conjunctive account of the type suggested by Bayne and Chalmers is natural. To be access conscious that p is to possess the ability to use the content p for the

¹¹ The phenomenology of proprioception is, likely, more complicated, but as I believe that it is fundamentally the same as external perception, I will assume here that it can be treated in the same manner.

immediate rational control of thought, speech, and action. Having these abilities in a unified way, or, to use perhaps more natural language, jointly having these abilities with respect to p and q is, on this view, naturally understood as having these abilities with respect to p and q . But consider some cases and intuitions in the vicinity. Consideration of these shows two things. First, which is relevant to this section, that the condition of p and q actually being able to be accessed on p and q being jointly accessible ought not to be accepted. If experiences are access conscious, and thus to have an experience is to possess an ability, then these abilities will behave like other abilities—of which the condition is not generally true. Second, which is relevant to later discussion in Sect. 6, that access unity is not ontologically introductory.

Consider owning a car and a bicycle. If one jointly owns a car and owns a bicycle, one owns a car and a bicycle. If one is jointly qualified to be a lawyer and qualified to be an accountant, then one is qualified to be a lawyer and an accountant. We should be careful of drawing certain conclusions from such considerations. To adapt a famous point from Ryle (1949), it would be quite wrong to say that there could be someone qualified to be a (lawyer and an accountant) who owns a (car and a bicycle). There are not such things if this means anything more than just someone qualified to be both a lawyer and an accountant or who owns both a car and a bicycle. There are not such things as (lawyer and accountant)s, where someone who is a lawyer and an accountant is also a (lawyer and accountant), or (car and bicycle)s, where someone who owns a car and a bicycle also owns a (car and a bicycle). Matters are perhaps less immediately obvious when it comes to the unity of consciousness as the unity of two states gives rise in some cases to phenomenology and also to a nexus of abilities. But the same general point does apply. If one jointly perceives a circle and a square, one does not perceive a (circle and a square). Likewise, there is no modification to one's perceptions of a circle and a square in the sense that one jointly-perceives a circle and a square, where this modification marks a distinction between the joint-perception and the individual perceptions. There is no more a further type of perception, nor a further perception of the same type, than there is a further type of object perceived, or further object of the same type perceived. The cumbersome notion of "joint perception" is really the same underlying notion as that which prompts the introduction of subsumption, a total conscious state, and so on. What we should here say is that there is someone who is both a lawyer and an accountant, who possesses both a bicycle and a car, and who is perceiving both a circle and a square. Such an individual possesses a nexus of abilities that they would not otherwise possess, were they not in joint possession of a car and a bicycle, legal and accounting qualifications, and not jointly—i.e. in a unified way—perceiving both a circle and a square. The same caution regarding the infelicitous introduction of the extra items applies just as much in the latter as in the former cases. So, it is generally true that the pattern from the joint possession of some properties A and B does not introduce a property of A and B . Abilities are properties like any other in this respect, and so presumably having the abilities to A and to B does not mean that one has the ability to A and B .

Consider now the following further point regarding abilities. The joint possession and joint exercising of abilities can come apart. One could jointly possess the ability to A and the ability to B without being able to jointly exercise those abilities. This

is relevant to the current discussion as while possessing, for example, the ability to report that p is a necessary condition for perceiving that p , being actually able to carry out the report is not a necessary condition on possessing the ability to report. Thus, even if one cannot actually carry out the report, this does not mean that one does not possess the ability to report. An interesting feature of the Sperling case is that it is the exercising of the ability to report that p that prevents one reporting that q and vice-versa. However, the exercising of one ability preventing the exercising of another is no different from something other than an ability doing so, and in such cases the subject can still be said to jointly possess the abilities.

Consider the following example. I am sitting in a room with a window and a door. I am in possession of the ability to open the door and I am in possession of the ability to open the window. I am in joint possession of both abilities: they are unified. I could open the window and I could open the door. But the window and the door are far enough from each other that I could not open both at the same time. So, while in one sense I could be said to be able to open the window and the door, in another sense I could only be said to be able to open the window or the door. Now consider my possession of these two abilities along side my lack of the ability to whistle. As I am unable whistle (sadly true), there is no sense in which I could whistle and open the window, or whistle and open the door. There is a trivial sense in which I could whistle or open the door, and whistle or open the window, as one disjunct is satisfied. But the underlying fact that I am in possession of the abilities to open the window and the door renders my ability to *open the window or the door* quite different from my abilities to *open the window or whistle* and *open the door or whistle*. The latter attributions are true only in the trivial logical sense related to disjunction introduction, whereas the former reflects the underlying fact that both disjuncts are true, even though the abilities described could not be exercised simultaneously. In one case, I jointly possess the abilities to A and to B , but circumstances dictate that I cannot exercise both abilities at the same time. In the other, I simply do not possess one of the abilities.¹² The circumstances which could preclude the exercising of the ability to A can include other abilities.¹³ The abilities conferred by access consciousness are no different in this respect from other abilities. And just as one can jointly possess two (or more) abilities, A and B , without being able to jointly exercise A and B , so one can jointly perceive that p and q , and hence have the unified abilities to

¹² The ability to A does not require the further ability A' to exercise the ability to A . The ability—or, rather, “ability”, as ordinary talk misleads here—to exercise an ability is not really a further ability that one possesses, only when as a matter of fact one can exercise the original ability. To take this line would be to set off on a regress.

¹³ One could perhaps attempt to argue that in this scenario were I to be in possession of a tool with which to reach both the handles of the door and window, or longer arms, or the window and the door were closer together, and so on, I would be able to exercise both abilities simultaneously, whereas this is not so in the Sperling case and thus there is a disanalogy. But the Sperling case, like all access bottlenecks, arises because of a limited information processing capacity. Were I to have this capacity extended, an extension that is in principle no different from an arm extension, I would be able to exercise the abilities to report that p and that q simultaneously and thus there is no disanalogy.

report that p and that q , without being able to jointly exercise the abilities and thus report that p and q . This is what happens in access bottleneck cases.¹⁴

An objection that will be raised to this, implied by Bayne and Chalmers, is that an account of access unity which allows for the failure of the expressibility of the access to p and q is not a substantive account. It amounts to no more than saying that the subject experiences that p and that q , when, intuitively, unity demands more—the ability to access p and q . I would resist this objection. Cases of perceptual spatial and object unity result in content that accounts for the phenomenology of unity. For example, the shape and colour of the letters on the Sperling grid are bound, and the letters are all spatially related to one another in perception. These are facts about content which goes further than saying merely that the subject perceives, for example, both the colour and shape of each letter. In being so bound, they are presented to consciousness as spatially and object unified. This is what explains why the subject perceives a grid of spatially related letters. They are thus also functionally related in that the experiences are located at the same functional point in the system. As a result, the subject does have access unity to p and to q in the sense that they possess both the abilities to report that p and that q : these are two abilities co-situated within the same functional nexus. The subject can report on both in the sense that they are jointly poised for access. These two abilities thus form part of a nexus of abilities that would be lacking were they not so unified. Were p and q not access unified and thus part of the same functional nexus, both of p and q would not be available for access. Both p and q being located at the same point in the functional nexus therefore accounts for the subject unity of p and q (the experiences of p and of q are experiences located at the same point in this functional nexus, or subject, as opposed to another) in terms of access consciousness as being located at this point in the functional nexus is to be access conscious, but does not require that the subject be able to report that p and q . It is a downstream feature of our information processing capabilities that we are unable to exercise both of these abilities at the same time. But, as we have seen, there are many cases of this, and one's ability to access the content p or q is akin to the ability to *open the door or the window* as opposed to the ability to *open the window or whistle* in that it reflects one's underlying access to both p and q .

However, it may be pressed that the unity of consciousness—including the functional unity of consciousness—presents a different phenomenon than mundane examples of unities of abilities, such as being able to open the window and the door. The subject has the ability to report that p and the ability to report that q , but the crucial question for the unity of consciousness concerns not possession of these two abilities, nor the expression of these abilities individually, but the *expression of both*

¹⁴ Following on from footnote 12, I think it similarly mistaken to hold that being able to express the abilities A and B together is a further ability over and above the abilities to A and to B . If I can open the door and also whistle, I do not possess three abilities—to *open the door*, to *whistle*, and to *be able to express the abilities to open the door and whistle*. If this were so, the abilities that one possesses would both quickly multiply, and the regress problem mentioned in footnote 12 would also reappear. This means that the objection that access unity in the bottleneck case should account for the ability to *express the abilities to report on the three rows jointly* would be mistaken for these reasons.

together. And as the subject fails to express these together in a functionally unified way, no robust theory of unity has been presented. No similar robust theory of unity is presented in the door and window case—but that is because no robust theory of unity is required in this case. The unity of consciousness is different from a mere coincidence of abilities, such as one finds in the door and window case. This is perhaps the most intuitive reaction to the argument, but I believe that the response given above is sufficient to disarm the intuition that it is not the mere unified possession of abilities that is paramount but the unified expression of the abilities.

An influential representational account of unity is Tye (1995) and Tye's terminology is useful here. In Tye's terminology, both p and q are simultaneously, i.e. jointly or in a unified way, 'poised' for the immediate use in rational control of thought, speech, and action. Were they not so poised, they would not be unified—but cases where they are so poised but cannot be simultaneously accessed are not such cases. And this is key. p and q are unified in that they are both poised at the same functional point in the system—what happens "down stream" does not affect this. To have an experience that p is for the content p to be suitably poised. Likewise for q . And when both experiences are jointly poised, then they are unified in that the functional unity of p and q is present within the system—they are simultaneously and equivalently poised within the same functional point of the system. It is a short step from this to the poise of p and q , and then to the condition on this that the expressibility of p and q is a condition on p and q being poised and thus on the unity of the experiences of p and of q . However, I hope to have made the case against this last step concerning expressibility. The access unity of the access consciousness that p and q , i.e. the fact that the two are simultaneously poised, is compatible with the actual inaccessibility of p and q , i.e. the thwarting of the simultaneous access of p and q due to factors which intervene to prevent the joint exercising of the performances of the ability to access both p and q . The subject jointly possesses the abilities to access p and q . Both of these content are poised at the same functional point within the system, and both are available for joint downstream use in coordinated rational control of thought, speech, and action. It is a feature of our processing capacities, that the expression of the ability to access one precludes the expression of the other. But, as we have seen, this is a general feature of abilities and it does not preclude a robust notion of joint access in terms of the functional point in the system at which both are poised.

To return to the metaphysical point about unity being ontologically introductory, even in a thin sense, if one wants to say here that p and q is poised, then one can put it in this way if one wishes, as it means that both p and q are equivalently poised. But p and q is no more poised over and above both p and q being simultaneously poised within the same system than are there are such things as (car and bicycles) or (lawyer and accountant).¹⁵

¹⁵ A reviewer observes that one could argue that the Sperling experiment shows that there is a limitation on parallel processing, and then extend this to a limitation on having certain joint experiences. One could understand this as the view that the subject can experience (the majority of the detail of) each row individually but that the limitation of parallel processing means that they cannot experience (the majority of the detail of) all rows together. One might cast this as a case of switching between the rows or attending to different rows. One thinks that one has an overall experience of all three rows to the same degree of detail that one has of the row that one is attending to at any one time, but this in fact not so. Thus there

One may further object, though, that an account of unity in these terms is too thin to get a grip on our intuitions. One could argue, for example, that a non-phenomenological account of unity cannot accommodate the intuition that unity is phenomenological. I would resist this argument. It is not clear what intuitions different people have, and which ought to be central to a theory of unity. Indeed, the account does accommodate the clear cases of phenomenology of unity, namely the perceptual cases of object and spatial unity. It is, to my mind, far from straightforward that the other cases of unity, as discussed in Sect. 3, really are cases in which there is a phenomenology of unity.

However, a deeper worry may be pressed. Perhaps one could accept that the access-bottleneck shows that if experiences being access conscious means that they are abilities, then the inference from the experience that p and the experience that q being unified to the content p and q being able to be accessed ought to be rejected. However, this does not show that this account of access consciousness is sufficient as an account of the unity of consciousness. To show that there could be an account of access unity which meets the bottleneck challenge is not to show that this is correct as an account of unity. Again, there is intuitive pull to this objection, but my response would be twofold. First, again, it is not clear which intuitions should be given the most weight, and I hope to have made a good case that on various points, the charge that the presented account is counterintuitive can be met. Second, my aim is to present an account of the unity of consciousness in terms of access unity which accommodates the clear cases of a phenomenology of unity that results from perceptual spatial and object unity, denies that there is a phenomenology to unity per se, and that can meet the objections that have been raised against it. The argument, therefore, only speaks indirectly against opposing views by presenting this account as one worthy of consideration.¹⁶

Footnote 15 (continued)

is no unified experience at any time with same level of detail for all three rows. One row is privileged, so to speak, and the other two rows have a corresponding reduction in detail. This is of course, an empirical question, but it is an ingenious response to the Sperling case. Were this to be correct, it would save the access unity account from the objection. However, even if it is not correct, the presented view can account for the overall experience of all three rows. On this view, it is the unified poise of the experiences that explains the unity of the overview experience of the grid. The limitation in parallel processing comes downstream, when the subject tries to simultaneously express the abilities thus jointly conferred.

¹⁶ Thanks to the three reviewers for pressing on this point. One reviewer comments that there may be split brain cases where the unity of consciousness breaks down but there is still some sense in which there is simultaneous accessibility. The central question about such a potential case is how the functional nexuses in which the two states are poised are related. Analyzing the unity of consciousness in terms of access consciousness is extremely flexible for handling such cases (see §5 and §6).

5 Breakdown of spatial unity?

A further objection to be considered is that perceptual spatial unity can break down but a phenomenology of unity remain and that there must therefore a notion of “phenomenal unity” that perceptual spatial unity cannot account for (the point can be extended to include objectual unity; my response likewise).

Versions of the objection has been raised by Dainton (2000), Tye (2003) and Bayne (2010), among others. A hypothetical case is adduced where a subject is engineered such that there are two separate perceptual streams which enter into a single agent’s stream of consciousness.¹⁷ Dainton (2000, p. 67), for example, considers a case where the subject’s brain is envatted and their auditory and visual input comes in through the sensory organs of a head which is detached and can be moved around far from a body through which they continue to receive tactile, proprioceptive, and kinaesthetic input. His scenario culminates with the subject navigating their body through the depths of the ocean while their head is strapped to a toy plane being flown through the sky. Such cases are intended to motivate an acceptance of the possibility of ‘phenomenal unity’ in the absence of any perceptual spatial unity. The subject’s streams of perceptions from the head and body are unified, but there is a radical spatial disconnection between the two.

Were one in the—fortunate or unfortunate depending on your appetite for excitement—position of Dainton’s hypothetical subject, it seems to me that Dainton’s view that there would be unity in this case is correct. One would have unified experiences of bodily sensations and unified visual and auditory experiences and that these two would be subject-unified. Dainton inquires as to whether or not there would be any sense in which the two streams of experience would be spatially unified with each other. While it is clear that no specific spatial relation would be represented between the locations of one’s head and body, there is the possibility that the fact that there is a spatial relation between the locations would be represented (we could helpfully think of this, perhaps, in terms of *de dicto* but not *de re* representation). Dainton concludes that it would seem as though there is a spatial relation between the two separated streams of experience, but that this is due to the ‘imaginary’ perceptual field that is present in our awareness of our own perceptual capacities. The way that I understand Dainton’s view on this point could be recast as the view that we know that even if our sensory input is split in such a radical way and there is no spatial content which binds the different sensory inputs together, there is still a spatial relation between the contents. Something along these lines seems to me to be correct and this starting point brings out well my response to such cases were they to be raised against the view I am proposing.

Despite their vividness, these cases are to my mind essentially no different in respect of the present discussion to mundane cases such as looking at the corner of

¹⁷ Bayne (2010: pp. 261–262) inquires as to whether there are empirical cases of spatial disunity in sufferers of heautoscopy who simultaneously retain their normal subjective point of view while also seeming to perceive themselves from an external perspective. The following remarks apply equally to empirical cases—these, or others if they are genuine cases—as to theoretical cases.

the office and imagining what a lamp would look like if placed there. In this case, we have two conscious experiences—the visual experience of the corner of the office and the imagining of the lamp in the corner—both of which have phenomenology and are clearly unified but where, it seems to me, there is no phenomenology to this unity. I would apply the same reasoning to cases such as Dainton's. Examples such as Dainton's do show that it is possible for one to have two groups of perceptual experiences which are spatially unified within themselves but not spatially unified with each other. If one's view were to be that there must be spatial content which holds between unified experiences, then these cases would be problematic. However, this is not the view that I am defending. The view that I am defending is that where there is a phenomenology of unity, there is perceptual spatial and objectual content in terms of which this phenomenology should be analyzed. In cases such as Dainton's, I would argue that there is no phenomenology of unity. While one's initial response to this claim may be to find it intuitively implausible, reflection shows, I believe, that one's response should be the same as to cases such as imagining the lamp in the corner. It would be deeply peculiar, and would no doubt clash in a discomfiting way with one's awareness of, for example, the possibility of acting across both perceived spaces. But we should not infer more than this strangeness from the case, as it is in principle no different *qua* the phenomenology of unity from the mundane case.

As I write this, I have one eye on the corner of my office and I am vividly imagining how the standing lamp from the other room would look if I placed it in the corner. No doubt these are unified experiences. They are even spatially unified in that the perceptual and imaginatory experience are linked by a spatial content. However, this content, like Dainton's 'imaginary' perceptual field, is not perceptual content. It is therefore unlike the spatial relation that I am perceiving between my desk and the window behind it. Rather, this spatial content is a cognitive content, part of my judgement or imagining of how things would be were the lamp to be placed in the corner that I happen to be looking at. Consider a failure of perceptual binding as a contrast case. For example, we might cast a case of agnosia in which a subject is able to describe the colour and the shape of an object but not categorize or act on the object as a discrete object with both that colour and shape as a failure of object binding. Were we able to turn such a subject's binding on, we would introduce extra content into their experience. In contrast, when I imagine the lamp in the corner of the office, I only add more content into my stream of consciousness. There is no analog to the previous type of binding in this case, no process by which the relation between one content p and another q in my stream of consciousness gives rise to more than just the sum of p and q . Dainton's case differs only in that it involves a general awareness of there being a spatial relation between the two different streams of perceptual content.

Perhaps, though, the following type of case could be more problematic. Consider some type of agnosia in which there is spatial/object unity within perception, and hence a phenomenology of this unity, but there is a breakdown in access unity. Or the converse case in which there is perceptual access unity of the type that goes with object and spatial unity between the experiences that p and q but no spatial/object unity reflected in the content and thus no phenomenology. Spatial and object unities

and the related access unities should not come apart from each other. Such cases are hard to assess because of the complexity of the context in which they would occur. A breakdown in access unity can come in degrees: unity of verbal abilities or rational abilities can be wholly or partially compromised, or the ability to act towards the content in one way but not another can be compromised, for example, while other abilities remain unified. The network of unified spatial and object contents could also break down in various ways, perhaps splintering perception into a mosaic of partly unified experiences, perhaps with breakdowns in transitivity of binding with respect to one object, and so on. How could we pull apart the particular phenomenological breakdown from the retained access unity and vice-versa? The complexity of this problem would require a detailed investigation into each case. I lack the space for this here, and so I will note that this is a potential problem but one which, for these reasons, I believe may not be as serious as it may first appear.

6 A Comparison with other accounts and the notion of a conscious field

To close, I will compare some aspects of other accounts of the unity of consciousness and extend my argument to cover the notion of a conscious field, before making some final brief remarks regarding some further questions about the unity of consciousness. In Sect. 2, I contrasted the presented theory with those from Hill (1991, 2014, 2018), Bennett and Hill (2014) and Masrour (2014). The theory is similar to both in that it rejects a unity relation which is phenomenologically manifest in experience, and which composes the unified experiences into a whole, or subsumes them. Access unity is not in any way ontologically introductory—even in the thinnest of senses. It differs from both, however, in that it specifies a single relation of unity, access unity, and from Masrour’s in that the unifying relation is not experienced.¹⁸ In analyzing unity in terms of access consciousness, the theory is a (first-order) functional-representational account of unity, and naturally, therefore, some who hold versions of this type of theory of consciousness, such as Dennett (1991) and Baars (1989), have proposed explanations of the unity of consciousness within this framework. Van Gulick (2013), for example, raises some objections in a similar vein to Bayne (2010), as does Baumann (2007) and Shoemaker’s (1996) account is widely discussed and influential.

Perhaps the most instructive comparison to the account offered in this paper, though, is Tye (2003). The comparison will help bring out the similarities and differences with other prominent accounts in the literature. Tye, in fact, shares with Bayne, Chalmers, and Dainton a key aspect of the central motivating commitment to

¹⁸ Although, in his discussion of the bottleneck objection, Hill (2018) comments that access unity is “plausibly” sufficient for the unity of consciousness but its necessity is an open question (2018: 8), and he proposes that the “gists” of experiences being extractable in bottleneck cases might allow an access-unity theory to answer the bottleneck challenge. Nevertheless, he goes on to advocate the multiple-relations theory of unity and therefore rejects, at least for the moment, an access-unity theory.

which I object (along with Hill and Masrour). I am in agreement with Tye’s restriction of the phenomenology of unity to perceptual experience, and in opposition to Bayne, for example, who takes it to cover all conscious states. I am in some further agreement with Tye regarding the way that unity, when it is phenomenological, holds together the contents of experience, not the experiences themselves (‘Phenomenal unity is a relation between qualities represented in experience, not between qualities of experiences.’ Tye 2003, p. 36). I suggest that otherwise separate contents are unified through spatial and object unity introducing discrete object and spatial content, but Tye, in general agreement on the underlying motivation with Bayne, Chalmers, and Dainton, holds that there is a relation between the contents in virtue of which ‘phenomenal unity is a matter of simultaneously experienced perceptual qualities entering into the same phenomenal content’ (Tye 2003, p. 36). On Tye’s view, this relation has the following consequence.

[P]henomenal unity goes with the closure of perceptual experience under conjunction with respect to the unified qualities. Thus [in a case where] the loudness of a sound is phenomenally unified for person *P* with the brightness of a flash of light, the statements [*P* has an experience of a loud sound] and [*P* has an experience of a bright flash] jointly entail [*P* has an experience of a loud sound and a bright flash]’ (Tye 2003, p. 37).

It is this relation, in any form, that has such a consequence, in any form, that the presented theory dispenses with. On this point, I am again in agreement with Hill and Masrour. However one characterises this relation—subsumption, the same phenomenal content, Dainton’s relation of co-consciousness that holds between unified experiences, etc.—I see no need for it as access-unity can provide an account of the unity of consciousness and its phenomenology without recourse to such a phenomenological whole which subsumes or contains, in any logical sense, the unified experiences.¹⁹

A further point of disagreement between Tye’s view and that which I am proposing concerns the priority of the unity over the unified. Even though, as Tye makes clear in his discussion surrounding the passages quoted above, his view is not intended to involve the type of double-counting discussed with respect to lawyers and accountants, and cars and bicycles discussed in Sect. 4, I do not believe that Tye is correct in taking it to have the consequence that he does with respect to the priority of the unified whole over the individual experiences. If the experiences that *p* and that *q* are access unified, all this means is that both are suitably poised within the same system. One need not draw any further consequences from unity than this. Despite disagreeing over the nature of the unified “whole”, subsumption as opposed to a single content, both Bayne and Tye do agree that there some such entity of unity. However, they further disagree over the priority of the unified and the unity. Bayne, along with Dainton and Lockwood, hold that individual experiences are logically prior to the unified whole, of which he offers a mereological account in

¹⁹ Hurley (1998) argues that on a functional-representation theory, the unified contents must be semantically coherent, i.e. lacking in contradiction, and so on. This would raise some problems, perhaps, for a view like Tye’s but Bayne (2010) argues persuasively that this condition is not necessary.

terms of subsumption. Tye, along with Searle (2000), holds the opposite view that the unified whole is logically prior, and that subjects undergo not individual unified experiences but a single experience with an extended single content. This debate has prompted quite some discussion and my treatment of it will therefore be perhaps too brief.²⁰ Insofar as I understand this debate, I side with Bayne against Tye for the reason that I reject the initial motivation both for Bayne's subsumption and Tye's single phenomenal content. But if one accepts an account of unity that does not have the double-counting implication, it seems to me as though not much can turn on this question. Compare two opposing views. One's unified experience is a network of independent experiences which are subject to some relation such that they are said to be unified. One's unified experience is a single entity, which unlike on the previous view does not have independent experiences as something akin to parts, but there is a principled, metaphysically important way of dividing this single experience into those very same parts. If there is anything to this debate, of which I am unsure, as Bayne (2005, pp. 498–499) points out, if one parcels out the single experience into smaller parts via a functional or phenomenological criterion, the difference between the two views seems to be of questionable importance. It becomes either a verbal dispute regarding how one counts or individuates experiences and/or a dispute about whether to privilege phenomenology or functional role in describing those aspects of conscious experience in which we are most interested.

As a final point, I would like to extend my argument to the notion of a conscious field. On the presented view, the phenomenology of unity is analysed in terms of the contents of perceptual object and spatial unity. Access unity itself introduces no phenomenology. This view is at odds with the prevailing view on which there is in some sense a conscious field. Here, for example, are Bayne and Bayne & Chalmers expressing the relation.

Phenomenal unity is often in the background in discussions of the 'stream' or 'field' of consciousness ... We can say that what it is for a pair of experiences to occur within a single phenomenal field *just is* for them to enjoy a conjoint phenomenality—for there to be something it is like for the subject in question not only to have both experiences but to have them *together* (Bayne 2010, p. 11).

(I)t seems plausible that all of my visual experiences are subsumed by a single encompassing state of consciousness, corresponding to my visual field. More generally, my visual and auditory experiences might all be subsumed by a single encompassing state of perceptual consciousness. And it does not seem unreasonable to suppose that there is a single encompassing state of consciousness that subsumes all of my experiences: perceptual, bodily, emotional, cognitive, and any others. We can think of this last encompassing state of consciousness, for a given subject, as the subject's *total* conscious state. When it exists, a subject's total conscious state might be thought of as the subject's conscious *field* (Bayne and Chalmers 2003, p. 28).

²⁰ For example, Dainton (2010), Masrour (2014) and Lee (2014).

If one wishes to describe the unity of consciousness in terms of a conscious field, then one can. However, it seems to me that the notion is essentially pleonastic. It adds no more to say that experiences occur in a conscious field than to say that they are unified. Imagine that you are trying to learn a magic trick. After some time, you are able to manipulate a deck of cards in order to perform the trick: you now have an ability that you did not have before. How are things for you different from the way that they were before? Well, you are now able to perform the trick. That is, there is one more thing on the list, in the bundle, or in the collection, of things that you are able to do. And as they are tied together in this way, depending on circumstances and how they interact with each other, you can combine performances of them in different ways. There is no pressure to invoke anything like a phenomenal or subsumptive unity which encompasses this ability and your existing abilities. There is no “field of abilities”. If I whistle while performing the magic trick, one need not introduce some field which contains both performances. In this example, I am just simultaneously engaged in two performances. We can apply this analogy to the notion of a conscious field and its content.

We can extend this line of argument to more specific conscious fields. Martin (1992), for example, argues that the visual field delimits the space of things seen, and this delimiting is itself something manifest in experience. We can think of this in terms of a boundary. The visual field is at least a boundary around the things seen, and this boundary is in some way manifest when things are seen within it: either it is literally something seen (doubtful), or it is manifest as a phenomenological awareness of the limit. On any reading, though, it is not more legitimate to introduce a conscious field as a boundary, enclosure, or container in which conscious experiences are had than it is to introduce a subsumptive phenomenal whole, a single phenomenal content, and so on. Consider a field of consciousness as a manifest boundary which delimits experience. This is perhaps most intuitive as a description of the visual field, but the idea seems to me to be misplaced even in this case. Just as there are no such things as (lawyer and accountant)s, so the end of a list is not another item on the list, or the boundary of a collection another item in the collection. The visual field is the collection of all of the things seen at one time, collected in virtue of their being seen at that time. The structure of the visual field is the network of object and spatial unities that holds between the different visual contents. The boundary of the visual field is not something that one sees or is visually aware of. It is the point at which things cease to be seen. Likewise for other perceptual fields, including the general “field of consciousness” which contains all of one’s experiences. Once we have accounted for the phenomenology of spatial and object unity and the further effects of unity in terms of access consciousness, we seem to have accounted for everything. We have no need for any unity over and above the functional nexus and the accompanying phenomenological consequences of the contents of perceptual unity, no need for any notion of a conscious field or its cognates.

Let me close by briefly mentioning some further related questions. As stated at the outset, I have made no comment about the modal strength of unity. A great deal turns on whether unity is necessary, in which circumstances it can break down (if possible), and what the consequences of this are. The account that I have given is neutral on these questions. If unity does break down, the view understands this as a

breakdown in access unity and a breakdown in access unity is quite a flexible tool with which to explain various psychological and behavioural disorders that may display plausible symptoms of disunity. Access unity can break down in various ways—in respect of verbal abilities, or capacities to act, for example—while being retained in other areas. If unity is necessary, then the account will face the usual problems with this account—seemingly two centres of consciousness in split-brain cases, for example. It seems no worse placed, and given the flexibility of access unity perhaps better placed, than other accounts in this respect. Another, related, issue that I have not touched on is the relation between the self and the unity of consciousness. Again, the view is neutral on the various options, although it fits most naturally with the understanding of the self as no more than the result of the coordination of functional-representational processing. Finally, the functional-representational account of unity that I have defended is a first-order account, on which conscious experience is a matter of the functional role of the processing of states with first-order content. The view, however, is equally capable of being understood in terms of a higher-order or reflexive theory of consciousness.

7 Conclusion

My focus in this paper has been on the phenomenology of the unity of consciousness. When two conscious experiences are unified, if there is a phenomenological consequence of this unity it is because the unity results in contents of perceptual spatial and object unity. There is no phenomenology *qua* unity in other cases. The further, non-phenomenological, effects of unity can be accounted for in terms of access unity. This view of unity can be defended both from the objection that as there are cases in which experiences of *p* and of *q* are unified but *p and q* is not accessible the account of access unity is insubstantial, and from the objection concerning the possibility of a breakdown in spatial unity. As a consequence, we need not introduce such notions as subsumptive or phenomenal unity, and we need not introduce any substantive notion of a conscious field.

Acknowledgements I would like to thank members of audiences at Hokkaido University, Huaqiao University, and Hiroshima University, in particular Tim Bayne, for their helpful comments and suggestions, and also Akiko Frischhut for detailed discussion and comments on previous drafts.

References

- Baars, B. J. (1989). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.
- Baumann, P. (2007). Experiencing things together: What's the problem? *Erkenntnis*, 66, 9–26.
- Bayne, T. (2005). Divided brains and unified phenomenology: A review essay on Michael Tye's consciousness and persons. *Philosophical Psychology*, 18(4), 495–512.
- Bayne, T. (2010). *The unity of consciousness*. Oxford: Oxford University Press.
- Bayne, T., & Chalmers, D. J. (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The unity of consciousness* (pp. 23–58). Oxford: Oxford University Press.

- Bennett, D. J., & Hill, C. S. (2014). A unity pluralist account of the unity of experience. In C. S. Hill & D. J. Bennett (Eds.), *Sensory integration and the unity of consciousness* (pp. 233–234). Cambridge, MA: MIT Press.
- Block, N. (1995). On a confusion about a function of consciousness. *Brain and Behavioral Sciences*, 18(2), 227–247.
- Dainton, B. (2000). *Stream of consciousness: Unity and continuity in conscious experience*. London: Routledge.
- Dainton, B. (2010). Phenomenal holism. *Royal Institute of Philosophy Supplement*, 67, 113–139.
- Dennett, D. (1991). *Consciousness explained*. New York: Little, Brown, and Co.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge, MA: MIT Press.
- Garson, J. (2001). (Dis)solving the binding problem. *Philosophical Psychology*, 14(4), 381–392.
- Hill, C. S. (1991). *Sensations: A defence of type materialism*. Cambridge: Cambridge University Press.
- Hill, C. S. (2014). Tim Bayne on the unity of consciousness. *Analysis*, 74(3), 409–509.
- Hill, C. S. (2018). Unity of consciousness. *Wiley Interdisciplinary Reviews: Cognitive Science*, 9(5), e1465.
- Hurley, S. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.
- Kriegel, U. (2009). *Subjective consciousness: A self-representational theory*. Oxford: Oxford University Press.
- Lee, G. (2014). Experiences and their parts. In C. Hill & D. Bennett (Eds.), *Sensory integration and the unity of consciousness* (pp. 287–322). Cambridge, MA: MIT Press.
- Lockwood, M. (1989). *Mind, brain and the quantum*. Oxford: Blackwell.
- Martin, M. (1992). Sight and Touch. In T. Crane (Ed.), *The contents of experience* (pp. 196–215). New York: Cambridge University Press.
- Masrour, F. (2014). Unity of consciousness. In defense of a Leibnizian view. In C. S. Hill & D. J. Bennett (Eds.), *Sensory integration and the unity of consciousness* (pp. 323–346). Cambridge, MA: MIT Press.
- Prinz, J. (2012). *The conscious brain: How attention engenders experience*. New York: Oxford University Press.
- Roelefs, L. (2014). The dimensions of the conscious field. *Journal of Consciousness Studies*, 21(7–8), 88–104.
- Rosenthal, D. M. (2005). *Consciousness and mind*. Oxford: Oxford University Press.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson.
- Searle, J. (2000). Consciousness. *Annual Review of Neuroscience*, 23, 557–578.
- Shoemaker, S. (1994). Phenomenal character. *Noûs*, 28(1), 21–38.
- Shoemaker, S. (1996). Unity of consciousness and consciousness of unity. In S. Shoemaker (Ed.), *The first-person perspective and other essays*. Cambridge: Cambridge University Press.
- Shoemaker, S. (2003). Consciousness and co-consciousness. In A. Cleeremans (Ed.), *The unity of consciousness* (pp. 59–71). Oxford: Oxford University Press.
- Tye, M. (1995). *Ten problems of consciousness*. Cambridge, MA: MIT Press.
- Tye, M. (2003). *Consciousness and persons: Unity and identity*. Cambridge, MA: MIT Press.
- Van Gulick, R. (2013). Phenomenal unity, representation and the self. *Philosophy and Phenomenological Research*, 86(1), 209–214.