Philos Stud (2016) 173:3199–3221
DOI 10.1007/s11098-016-0658-7

The unity of consciousness, within subjects and between subjects

Luke Roelofs[1]

Published online: 16 March 2016

Abstract The unity of consciousness has so far been studied only as a relation holding among the many experiences of a single subject. I investigate whether this relation could hold between the experiences of distinct subjects, considering three major arguments against the possibility of such 'between-subjects unity'. The first argument, based on the popular idea that unity implies subsumption by a composite experience, can be deflected by allowing for limited forms of 'experience-sharing', in which the same token experience belongs to more than one subject. The second argument, based on the phenomenological claim that unified experiences have interdependent phenomenal characters, I show to rest on an equivocation. Finally, the third argument accuses between-subjects unity of being unimaginable, or more broadly a formal possibility corresponding to nothing we can make sense of. I argue that the familiar experience of perceptual co-presentation gives us an adequate phenomenological grasp on what between-subjects unity might be like.

Keywords Unity of consciousness · Consciousness · Philosophy of mind · Phenomenology · Metaphysics · Experiences · Perception · Mereology

1 Introducing between-subjects unity

Much recent discussion of consciousness has focused on its unity, the way that our many experiences are somehow had together, forming a single integrated conscious field. When I see the screen in front of me, hear the sound of traffic, and reflect on

✉ Luke Roelofs
luke.roelofs@utoronto.ca; luke.roelofs@anu.edu.au;
https://researchers.anu.edu.au/researchers/roelofs-lmf

[1] School of Philosophy, Research School of Social Sciences, College of Arts and Social Sciences, The Australian National University, Canberra, ACT 0200, Australia

my memories from yesterday, there is not only 'something it is like' for me to experience each of these things, but also 'something it is like' for me to experience all of them together. This contrasts with three similar experiences that occur separately, one to me, one to you, and one to a third individual somewhere else. More precisely, we can draw two contrasts here: first, in the former but not the latter case, all three experiences are 'subject-unified', meaning simply that they belong to the same subject; second, in the former but not the latter case, the experiences are 'phenomenally unified', standing in some sort of experiential relation such that they together form an experiential whole.

The first fact concerns simply the 'ownership' of the experiences, their relation to a subject; the second concerns their experiential character (individual and collective). The driving idea of the literature on the unity of consciousness is that we can discuss and investigate this second phenomenon—the relatedness of experiences as such—in its own right, not merely as a derivative shorthand for how experiences relate to a subject. As a result, we can frame and evaluate various different theses about their relations—a prime example being the following thesis (taken from Bayne and Chalmers 2003, p. 24):

> **Unity Thesis:** Necessarily, any set of conscious states of a subject at a time is unified.

If 'unified' here meant 'subject-unified', this thesis would be trivial: but with 'unified' taken to mean 'phenomenally unified', it becomes a substantive and interesting claim—the claim that subject-unity implies phenomenal unity (Bayne and Chalmers 2003, pp. 24–27; Shoemaker 2003, pp. 59–65; Tye 2003, pp. 133–165; Bayne 2010; Schechter 2010). Moreover, once phenomenal unity is distinguished from subject-unity, we can even frame and evaluate purported *analyses* of subjects and their identities in terms of phenomenal unity (e.g. Dainton 2008; Shoemaker 2003, pp. 66–71; Bayne 2010, pp. 269–294).

However, there has not been any similar discussion of the converse thesis, that phenomenal unity implies subject-unity. Although theorists have debated the possibility of *disunified subjects*, breakdowns of phenomenal unity in the experience of a single subject, they have entirely neglected the possibility of *between-subjects unity*, phenomenal connections between the experiences of distinct subjects. The reason, I think, is that the impossibility of between-subjects unity is thought too obvious to merit discussion: after all, how could the very same relation which knits together my experiences, also connect my experiences with someone else's experiences, which I do not have?

In this paper I will defend the possibility of between-subjects unity, arguing that once we recognise that phenomenal unity is a real relation distinct from subject-unity, there is no compelling reason why it could not, under the right circumstances, connect the experiences of distinct subjects with one another. I will consider three major lines of argument that might seem to rule this out, which I call the argument from subsumption, the argument from interdependence, and the argument from unimaginability. As well as a defending the bare possibility of between-subject unity, I will gradually develop a more fleshed-out positive account of how between-subjects unity might work, elaborating it further in the course of responding to each

of these three objections. One particularly significant qualification, which emerges in Sect. 3, is that between-subjects unity is probably only possible in cases where subjects mereologically overlap, either sharing parts of being parts of one another. What emerges is a coherent and intelligible, but very unfamiliar, conception of experiences and subjects, on which properly-related conscious subjects can intelligibly constitute a single overarching consciousness without themselves being dissolved into the whole.

This investigation can throw fresh light on the vexed question of 'phenomenal holism': of whether a total conscious state is 'built up out of' its parts, of if they are rather 'mere abstractions from' it (see Gurwitsch 1964; Sprigge 1983; Searle 2000; Bayne 2010, pp. 225–249; Dainton 2010; Chudnoff 2013; Koksvik 2014). One of the persistent difficulties in this debate is to get clear on what exactly is at issue when we ask whether the whole or its parts are 'prior', and it is here that the question of between-subjects unity becomes useful. If between-subjects unity were impossible, we could say that unified experience is 'holistic' in at least this sense: its component experiences are not things that any subject could experience just by themselves,[1] given their actual relations to each other: to be experienced by themselves they would need to lose their unified character And if, as I argue, between-subjects unity is possible, then we can conversely say that unified experience is 'atomistic' in at least this sense: its components are the sort of thing that could coherently characterise a subject's experience all by themselves, even given their unity relations to other components. Of course this will not resolve the general question of phenomenal holism (not least because there remains the question of how *this* criterion of holism/atomism relates to the other criteria that have been advanced), but it provides a useful way to clarify what is or is not at issue.

As well as phenomenal holism, the question of between-subjects unity is intimately connected to a very deeply-rooted idea about the mind: that minds are in some sense both both *mereologically simple*, not composed of other minds, and *metaphysically isolated*, not overlapping with or sharing experiences with other minds. This is the idea at work when Descartes, Plotinus, and myriad others[2] argue that the mind must be immaterial, because any material thing would be composite. This is the topic of Kant's second paralogism (1998, A351–A354), the source of Nagel's intractable difficulties in making sense of the split-brain patient (1971), the reason why Putnam (1965/2003, pp. 215–216) and Tononi (2012, pp. 59–68) impose an anti-nesting requirement or 'exclusivity principle' on their accounts of consciousness. James expresses this idea when he says "The breaches between… thoughts are the most absolute breaches in nature" (1890, p. 226), Coleman when he says that "our notion of a mind, like our notion of a subject, is precisely the notion

[1] I will speak of subjects 'experiencing experiences' interchangeably with speaking of them 'having experiences' or 'undergoing experiences': this is merely stylistic, and not meant to carry any theoretical implications.

[2] See Plotinus (1956, pp. 255–258, 342–356), Proclus (1963, p. 163), Avicenna (1952, pp. 47ff), Descartes (1985, Volume 2, p. 59), Butler (1860), Mendelssohn (2002), Clarke (1978, Volume 3, p. 759), Bayle (1991, pp. 128–134), Lotze (1894, p. 158) and Brentano (1987, pp. 290–301). For a review, see Mijuskovic (1984), Schachter (2002), and the articles in Lennon and Stainton (2008).

of a discrete, essentially inviolable sphere of conscious–experiential goings-on." (2012, p. 145), and Zuboff when he says of a brain-splitting thought experiment that although "we may easily think of the brains themselves in terms of fractions… [and] fall back on talking about there being half of the original brain with you and half now over there with the other. But one could never talk about the subject or his experience like that." (1990, p. 41) Barnett (2008) even argues that it is this intuitive conception of the simple mind that lies behind certain influential thought-experiments, such as Searle's 'Chinese Room' and Block's 'Homunculus-Head' (1978).

I will call this resistance to viewing minds as composite the 'Anti-Combination' intuition (cf. Roelofs 2015). The influence and depth of Anti-Combination has not, I think, been adequately recognised. Ultimately I think it is a false intuition, but it is not something that we can simply dismiss: it must be carefully understood and systematically dismantled. Obviously that task goes beyond this paper, but defending the possibility of between-subjects unity is a crucial part of it. Supporters of Anti-Combination have often employed arguments along the following lines: the unity of consciousness could not be accounted for by any system comprising multiple distinct subjects, however they might be related. As Kant puts it, "representations that are divided among different beings… never constitute a whole thought" (Kant 1998, A353). In affirming the possibility of between-subjects unity, we deny this premise: there *is* a relation among subjects that can account for the unity of the consciousness of the composite they form, namely the relation of each having experiences which are unified (despite the distinctness of the subjects) with those of the other.

If I am right that between-subjects unity is possible, that opens up major new options in understanding several current debates in the metaphysics of mind, such as those about the possibility of group consciousness (Block 1978; Knobe and Prinz 2008; Huebner 2011; Schwitzgebel 2014), the split-brain syndrome (Nagel 1971; Tye 2003, pp. 109–132; Bayne 2010, pp. 189–221), hypothetical person-splitting (Parfit 1971; Lewis 1976; Zuboff 1990), and the ontological status of overlapping large parts of conscious beings (Merricks 2001, pp. 95–111; Unger 2006, pp. 377ff). In particular, panpsychists about consciousness (Seager 1995; Strawson 2006, Chalmers 2015), who hope to explain human consciousness by postulating primitive consciousness in all matter, need to explain how the consciousness of atoms and molecules could constitute that of a human being (Goff 2006; Basile 2010; Roelofs 2014; Chalmers Forthcoming).

What connects these debates is that they involve conscious subjects which are in some intuitive sense composed of other conscious subjects, whether that is a group agent composed of many individual members, a split-brain patient whose hemispheres are regarded as distinct subjects, or a normal human as explained by panpsychism. Different views of the ontology of subjects yield different ideas of what this means, and what 'part' of a subject is: if subjects are physical things like organisms or organs, then their parts will be sections of those physical things, while if subjects are (say) constituted by a stream of experiences, then their parts might be any subjects constituted by subsets of their experiences. But however we spell out the idea of subjects composing subjects, allowing for between-subjects unity

provides a new way to explain within-subject unity, or a lack thereof, in a composite being's experiences. If we can make sense of between-subjects unity among the experiences of the conscious parts involved, this relation could explain an/or constitute the within-subject unity of the whole. Perhaps a social group would have unified consciousness iff its members enjoyed between-subjects unity; perhaps human beings have unified consciousness in virtue of their hemispheres or neurones or atoms enjoying between-subjects unity (and lose it in the split-brain case because their parts cease to stand in that relation). If between-subjects unity is impossible, however, this possibility is closed down and unified consciousness in the composite will always be in some sense emergent.

In Sect. 2 I outline the three arguments against between-subjects unity that I will consider; Sects. 3, 4, and 5 then address them in turn. In Sect. 3 I show that between-subjects unity is compatible with the attractive 'subsumption' account of phenomenal unity, as long as it allows for the possibility of limited 'experience-sharing', a possibility which I defend against two major objections, and which implies a restriction of between-subjects unity to cases where subjects mereologically overlap. In Sect. 4 I show that between-subjects unity, and experience-sharing, are compatible with extensive, even all-pervasive, phenomenal interdependence among experiences, and that extant arguments against this compatibility rest on an equivocation. Finally, in Sect. 5 I consider whether the defender of between-subjects unity can provide a positive characterisation of its phenomenology, arguing for a characterisation based on the phenomenon of perceptual co-presentation, our awareness of objects as present but not given.

2 Three arguments against the possibility of between-subjects unity

Between-subjects unity is a very peculiar idea at first sight, but can we actually articulate any reason to rule it out? Of course, if we initially pick out 'the unity of consciousness' by means of a contrast between experiences within one mind, and experiences spread across other minds, it is hard to say that it in fact holds between distinct minds *in general*, without losing our grip on the phenomenon. But that is far from showing that this relation can *never* obtain between experiences in distinct minds. And insofar as the burden of proof lies with the defender of a necessity claim rather than the defender of the converse possibility claim, we should not rule out between-subjects unity without some positive argument. Here are three attempts to provide such an argument, in descending order of precision.

The Argument from Subsumption An essential feature of phenomenal unity is 'subsumption': two experiences which are unified are parts of a subsuming experience which comprises both. But an experience needs a subject, and so this subsuming experience must belong to some subject, who therefore experiences both of the component experiences as parts of this whole. Therefore any case of phenomenal unity must involve a single subject having both of the unified experiences, and since experiences cannot be shared, it is the only subject having either. Therefore between-subjects unity is impossible.

The Argument from Interdependence An essential feature of phenomenal unity is the dependence of each experience's character on the others': when many experiences are unified, they usually, if not always, each *feel different* from how they would feel in isolation. But if a subject were experiencing one but not the others, they would experience it in isolation. So that experience would somehow be both dependent on and independent of the others, which is impossible. Therefore between-subjects unity is impossible.

The Argument from Unimaginability Even if the above arguments were to fail, all that would show is that between-subjects unity is a *formal* possibility, not that any real possibility actually corresponds to this theorist's invention. And it seems impossible to form any positive, substantive, conception of what between-subjects unity would be like: in particular, we cannot imagine what it would be like to be one of the subjects involved, experiencing the unity from one side only. What could it be like for me, to have one of my experiences unified with another experience that I did not have? In light of this unimaginability, it is more reasonable to think that the whole idea of between-subjects unity is misguided.

In the next three sections, I consider these three arguments in turn, seeking not only to defeat them but also to thereby elaborate a positive account of between-subjects unity by observing what such an account would have to say in order to resist these attacks. None of these arguments is likely to seem compelling to all readers: in particular, the first will not move readers who do not regard subsumption as a necessary feature of unity, the second will not move those who do not accept the phenomenological claim of pervasive phenomenal interdependence, and the third will not move those who see no need for theoretical posits to be positively imaginable by us. With those readers, therefore, my defensive task-removing objections to between-subjects unity—is much easier. Yet my constructive task-developing a positive account of between-subjects unity—is harder, since they do not recognise all of the constraints I am trying to make this account fit. I invite them, therefore, to view my positive account as just one option among many, and to formulate and defend other accounts of between-subjects unity. But in this paper, I will try to take on board the driving intuitions behind all three of the above arguments, so as to show that even then, between-subjects unity is possible.

3 The argument from subsumption

A very natural and appealing way to think of phenomenal unity is that when two experiences are unified, there is not only 'something it is like' to have one, and 'something it is like' to have the other, but also 'something it is like' to have both together. If there is such a 'conjoint phenomenal character', it seems there must be an experience whose character it is, an experience which both of the unified experiences are in this sense 'parts of'.

The popular notion of experiential 'subsumption' aims to capture this idea by saying that whenever two experiences are unified, they are subsumed by another experience, which contains them as parts in at least the sense that by having that experience a subject would automatically have the other two. The 'subsumptive

account' of phenomenal unity (defended at length by Bayne 2010; cf. Bayne and Chalmers 2003) takes this to provide an *analysis* of phenomenal unity, but even theorists who reject such an analysis may accept that wherever there is phenomenal unity, there is subsumption.[3]

So we can take, as the first premise of our first argument against between-subjects unity, the claim that whenever there are two experiences that are unified, there is also a third experience that subsumes them both. But we will need additional premises, starting with the following (taken from Dainton 2011, p. 246):

> **Exclusivity Principle:** if an experience e_1 belongs to a subject S_1, it belongs ONLY to S_1, so e_1 cannot also (and simultaneously) belong to a distinct subject S_2.

Denying the exclusivity principle means allowing for experience-sharing: a single token experience simultaneously belonging to two distinct subjects. For instance, we might suppose that both a whole human being and that human being's brain are having experiences, but that it is the very same token experiences that each is having. Or we might suppose that one part or subsystem of the brain is sharing some but not all of the experiences of the whole tfabrain (e.g. mfy left hemisphere and I share tactile experiences from my right hand, but not from my left). The particular forms of experience-sharing that might occur are not my topic here, but just the bare metaphysical possibility.

The exclusivity principle has generally been regarded as plausible, if not self-evident (Cf. James 1890, p. 226; Dainton calls it plausible), but I will shortly argue against accepting it in the unqualified form Dainton gives. For now I will focus on showing how it supports an argument against between-subjects unity. The final premise of that argument is the following:[4]

> **Ownership Principle:** For any experience, there must be a subject who experiences it.

If we affirm *both* the exclusivity principle and the ownership principle, then the subsumptive account directly rules out between-subjects unity, as follows:

Suppose for *reductio* that we have two distinct subjects, S_1 and S_2, who have distinct experiences e_1 and e_2, which are unified with one another.

1. Since e_1 and e_2 are unified, the subsumptive principle implies that there is an experience e_3 which subsumes them.
2. The ownership principle implies that there must be a subject for e_3: call this subject S_3.
3. The definition of subsumption implies that S_3 experiences both e_1 and e_2, for it experiences e_3, which subsumes them.

[3] In particular, adherents of the 'one-experience' account, on which phenomenal unity is analysed as there being only one experience per subject at a given moment, are committed to subsumption being a necessary condition of unity, since subsumption is reflexive just as (improper) parthood is (Cf. Bayne 2010, pp. 29–30).

[4] For a recent defence of this principle, see Strawson (2013); for a recent attack, see Coleman (2012). For a deflationary account of its meaning which nevertheless makes it a necessary truth, see Parfit (1999).

4. But the exclusivity principle implies that there can be only one subject for each of e_1 and e_2, and thus that S_3 must be identical to both S_1 and S_2.
5. But then by the transitivity of identity, S_1 and S_2 are identical, contra the initial supposition that they were distinct.

There are broadly two ways to deflect the argument from subsumption. The first way is to deny either the subsumptive principle, or the ownership principle, thereby allowing for cases of between-subjects unity that give rise to no third subject. If experiences do not require subjects, then there need not be any subject for e_3 in the above example: two subjects might experience unity between their experiences just if there is a third experience subsuming both, which occurs but belongs to no subject. And if unity does not entail subsumption, there need not be any e_3 at all. The second way is to deny the exclusivity principle, allowing for experience-sharing. Then S_3 can be distinct from one or both of S_1 and S_2, despite sharing with them e_1 and e_2. Note that this second response, by accepting subsumption and ownership, implies that all cases of between-subjects unity will also involve within-subject unity, since at least one subject must have the subsuming experience e_3 and thus have both e_1 and e_2; yet this does not remove between-subjects unity, since S_1 and S_2 may still experience the unity relation from one side only.

I think there are two reasons to prefer the latter of these two options (retaining subsumption and ownership while denying exclusivity). Firstly, the subsumption principle is a very plausible necessary condition on unity (even if we remain neutral about subsumption as an analysis of unity). If we said that there could be unified experiences where there was nothing it was like to experience them both together—no conjoint phenomenal character—then it would be hard to see why that relation deserved to be called unity. Similarly, the ownership principle is plausibly just a corollary of what 'experience' means, at least as long as we are working with a suitably 'thin' notion of subject. Experiences are characterised by 'what it is like' to undergo them—and this seems to imply a reference to their undergoing, and thus to something which undergoes.[5]

Secondly, denying subsumption or ownership but keeping exclusivity would imply that while there might be between-subjects unity, it will never play any role in explaining the consciousness of a composite subject—for the subsuming experience will be 'unowned'. In particular, the within-subject unity of our own human minds could not be explained by or even interestingly related to any between-subjects unity of the parts of us, on pain of denying that we are subjects. So if we are interested in between-subjects unity in part for its potential to undermine what I earlier called the Anti-Combination intuition, we are driven to deny exclusivity.

Fortunately, I think that the exclusivity principle can reasonably be denied, or at least qualified. The idea of that principle is that no subject can share its experiences with 'another' subject, but there are two ways to read this 'another': we might mean

[5] By contrast, if we think of 'subject' as implying some rich and complex structure, like a persisting stable psychology, then it will be less plausible that every experience must have such a subject, but perhaps more plausible that there could be only one such subject for any experience. Hence it is important that I am not using 'subject' in that sense.

simply a 'distinct', i.e. non-identical, subject (as in Dainton's formulation), or we might mean 'discrete', i.e. non-overlapping, subjects–subjects which are not parts of one another, and do not have any parts in common.[6] That yields the following four theses:

Strong Exclusivity (SE): A single experience cannot belong to multiple distinct subjects.
Weak Exclusivity (WE): A single experience cannot belong to multiple discrete subjects.
Weak Sharing (WS): A single experience may belong to multiple distinct subjects.
Strong Sharing (SS): A single experience may belong to multiple discrete subjects.

WS is the negation of SE, while SS is the negation of WE; SE entails WE and SS entails WS. But, crucially, WE and WS are compatible—experiences might be weakly shareable yet also weakly exclusive. There are two obvious arguments for exclusivity, one formal and metaphysical, appealing to individuation criteria, one intuitive and epistemological, appealing to the idea of 'privacy'. But both fail as defences of Strong Exclusivity, though they may succeed as defences of Weak Exclusivity.

The first argument for exclusivity is that it simply falls out of the most natural way to individuate experiences, namely as instantiations of experiential properties.[7] Property instantiations, plausibly, are individuated by three factors: the property instantiated, the object that instantiates it, and the time at which the latter instantiates the former. But the objects that instantiate experiential properties are subjects, and so experiences are, it seems, individuated by their subjects. But then any experiences of distinct subjects must necessarily be distinct, since they differ in one of the dimensions that individuate them (Cf. Bayne 2010, pp. 24–29 for an argument that this is the right way to individuate experiences).

The problem with this argument is that it ignores the way that this rule for individuating property instances must be adjusted to accommodate cases of composition and overlap. It is intuitively plausible that token physical properties are often shared between parts and wholes. For instance, an object with an uneven surface seems to share that instance of unevenness with its surface (and all sections of it that include that surface). A blue surface seems to share, with each of its blue

[6] As noted in Sect. 1, we can spell out the meaning of 'discrete' in different ways for different views of what sort of thing subjects are. I believe Weak Exclusivity is equally defensible on most natural ways of spelling this out, but to do justice to this claim here would require elaborating a complete account of how subjects can be composed of subjects, which is a large task. I do attempt this larger task in other work, but my focus in this paper is on whether the nature of the unity relation itself proves an obstacle; I maintain that it does not.

[7] I define 'experiential properties' as properties which are individuated by what it is like to have them. This is narrower than simply 'properties which it is like something to have': if I can desire X both consciously and unconsciously, then there is sometimes something it is like to have the property 'desires X', but what it is like is not essential to that property. Cf. Bayne and Chalmers (2003, pp. 30–31), on 'phenomenal states' versus 'phenomenally conscious mental states'.

subsections, their particular instances of blueness. A car may be scratched when its roof is scratched, and this seems to involve only one instance of the property 'being scratched': similarly for being perforated, or wounded. If I can be said to be 'metabolising alcohol' when my liver is, it seems wrong to think of this as two instances of that activity. In all of these cases it seems natural to say that there is a single instance of the property in question, which can be truly ascribed both the whole and the part. Something similar goes for cases of overlap: if conjoined twins share a liver, both are metabolising alcohol but there is still only one instance of that activity, not two. And likewise with 'material constitution', the relation an object bears to the 'stuff' that makes it up, which we may analyse as complete overlap: a state which instantiates a certain shape-property might be distinguished from the hunk of metal that constitutes it, but this seems to involve only one instance of that shape-property, not two (Cf. Baker 1999).

These cases are sufficiently different from cases where there are two separate and independent instances of the same property (two blue things side-by-side, two distinct scratches or wounds, etc.) that we can reasonably consider it a requirement on any theory of properties that it be able to account for these instances of token-sharing. Moreover, refusing to recognise token-sharing leads to awkward forms of 'double-counting': if each of my limbs, torso, and head weighs 10 kg, and I weigh 60 kg, we would risk concluding that the scale will read '120 kg' when I and my body parts are on it.

So I conclude that while it is a good idea in general to individuate property instances by their bearers, this must be adjusted somehow to accommodate the fact that wholes share property instances with their parts. This might be done by individuating instances by, for instances, a set of bearers (and requiring that no members of this set be discrete from each other), or by a 'minimal' or 'maximal' bearer (the bearer which does not have other bearers as parts, or is not contained in other bearers). Consequently, the argument against experience-sharing will fail in cases where the subjects sharing an experience are mereologically related rather than discrete. If my whole brain is a subject, for instance, it might not be able to share experiences with *other*, *discrete*, brains, but that would not rule out its sharing experiences with me, or with its own subsections, lobes, or hemispheres, if any of them were subjects.

The second argument I will consider against experience-sharing starts from the idea that conscious experience is essentially 'subjective' rather than 'objective', in some sense, and that this entails that experience is 'private', knowable directly only by one subject. This distinguishes it from the world of matter, all facts about which are 'public', equally knowable to all subjects. Of course the claim of privacy (we might also say 'privileged access') is only as well-defined as the idea of knowing something 'directly' rather than 'indirectly', but it will not matter here what account to give of this distinction. All that we need is for the notion of 'direct' knowing to connect privacy with exclusivity, via the plausible idea that *having* an experience is both necessary and sufficient for being in a position to know it directly.[8] Given this

[8] Indeed, the theses of privacy and exclusivity are sometimes not distinguished at all; for instance, Unger (1990) describes the thesis of 'the privacy of experience' thus: "Except for that particular subject himself, nobody else and nothing else can *have* that conscious experience that he has. As another gloss on this

link between *having* an experience, and being in a position to *know* it directly, accepting experience-sharing threatens to violate privacy, making experience public by allowing multiple subjects to know the same experience directly.

While the defender of experience-sharing might simply bite the bullet and deny privacy, I think a more conciliatory approach is to move to Weak Exclusivity, and maintain that this still vindicates the intuition of privacy. To see how, observe that as well as distinguishing Strong Exclusivity from Weak Exclusivity, we might distinguish the following two theses:

Strong Privacy: A single experience cannot be directly known by multiple distinct subjects.
Weak Privacy: A single experience cannot be directly known by multiple discrete subjects.

I think that Weak Privacy captures the intuitive force of privacy just as well as Strong Privacy, and Weak Privacy is compatible with experience-sharing among mereologically related subjects. In particular, Weak Privacy still contrasts sharply with the publicity of physical fact; to parallel the distinction between Weak Sharing and Strong Sharing, we can distinguish Weak Publicity from Strong Publicity, as the negations of Strong and Weak Privacy:

Weak Publicity: A single experience may be directly known by multiple distinct subjects.
Strong Publicity: A single experience may be directly known by multiple discrete subjects.

Physical facts are *strongly* public, while experiential facts are merely weakly public. Whereas a physical fact can in principle be known equally well by any subject, an experience is directly knowable only by those intimately and directly involved in its occurrence. The fact that more than one distinct subject may be intimately and directly involved in a single experience's occurrence need not render this contrast less significant.

Suppose that Weak Privacy, Weak Exclusivity, and the principles of Subsumption and Ownership, are all true: there could then be between-subjects unity only in cases where the unified experiences belong not only to two distinct subjects but also to a composite subject of which they are both parts. Note that this does not mean that the first two subjects must be parts of one another, or overlap at all: they might be entirely discrete subjects, but both parts of a third.[9]

Footnote 8 continued
idea, nobody else, and nothing else, can be *directly conscious of* the experience of that particular subject" (p. 40, emphasis added).

[9] For instance, if by some futuristic contrivance I was able to unify some of my experiences with some of another person's (perhaps using radio-transmitters implanted into our brains), it would not follow that we were no longer discrete subjects: what would follow was that the pair we formed, or some larger group which we both belonged to, must be or become a conscious subject.

4 The argument from phenomenal interdependence

The second argument against between-subjects unity starts from the idea that phenomenal unity between two experiences is not a purely external relation, but actually changes the intrinsic phenomenal character of each. Phenomenal interdependence seems most plausible in cases of gestalt perception, such as when three black circles with sections missing are perceived as occluded by a white triangle, but only if they are seen arranged in the right way with two other such circles. But others have argued for much more extensive everyday forms of interdependence, such as the effect of mood and context on the quality of sensation; some have maintained that such interdependence is all-pervasive among unified experiences (see e.g., Gurwitsch 1964, p. 120ff; Dainton 2000, pp. 181–213; Basile 2010, p.110; Chudnoff 2013). Dainton (2010, pp. 133–139) advances a slightly different argument for pervasive interdependence, arguing that the mere unification of an experience with others should be counted as a feature of its phenomenal character.[10]

Suppose that we accept thorough-going phenomenal interdependence; what does that have to do with between-subjects unity? Two recent papers (Basile 2010; Coleman 2013) have argued that, given phenomenal interdependence, experience-sharing is impossible, because if a single experience were part of two different total phenomenal fields, its interdependence with the other members of those fields would require it to have two different phenomenal characters. Given the arguments of the last section, an argument against experience-sharing is already a threat to between-subjects unity; indeed, both arguments are presented as posing a problem for constitutive panpsychism, precisely because they take that doctrine to require building up unified minds out of component minds somehow related, which they take to require experience-sharing. But we can also envisage a more direct argument against between-subjects unity based on phenomenal interdependence: in essence it would claim that between-subjects unity would require a single experience to have two different phenomenal characters, one based on the other experiences of its subject, the other reflecting also the experiences of another subject which it is unified with. I think that careful analysis will show that this argument, as well as Basile's and Coleman's, are invalid, but it is not straightforward to see how.

[10] The claim of thorough-going phenomenal interdependence sometimes prompts the following line of thought: if each 'component experience' depends for its character on the others, does it really make sense to regard them as real and distinct things? Or should we not rather treat them as mere abstractions from the total experience (cf. Dainton 2010; Chudnoff 2013; Koksvik 2014)? This can lead towards a 'one-experience' account of unity, on which our experiences are phenomenally unified just in that, strictly speaking, we each have only one experience at a time: the diversity we find in consciousness is merely diversity in the content of this experience. On this account (defended most notably by Tye 2003; cf. James 1890), the unity relation is simply identity: two co-conscious experiences are a single experience, perhaps picked out by different descriptions. Had my present total experience been different in any way, it would be a distinct total experience, and rather than saying that one of my experiences would have been different but others the same, we should say that my experience would have been a different one that resembled my actual experience in some respects. The one-experience account obviously makes it very difficult to make sense of between-subjects unity, so I am committed to rejecting it; fortunately, the arguments developed in this section help to show that phenomenal interdependence does not necessitate the one-experience view.

In this section I will first lay out Basile's and Coleman's arguments against experience-sharing, and show how an analogous argument might be constructed against between-subjects unity; then I will identify the common flaw in these arguments.

Basile bases his argument on two principles which he attributes to William James, and which together seem to pose a problem for experience-sharing:

> PHENOMENAL ESSENTIALISM: … for an experience, to be is to feel a certain way… in the case of experience, 'appearance' and 'reality' are one and the same.
>
> PHENOMENAL HOLISM: …within a person's total psychical whole, the nature of a single identifiable experience… is essentially determined by the other experiences occurring alongside it… within the whole… (Basile 2010, p. 107)

Phenomenal essentialism implies that experiences cannot be numerically the same while feeling different, but Phenomenal Holism implies that an experience will feel different in when unified with different sets of other experiences. Hence an experience cannot simultaneously be part of two different unified sets of experiences, and thus cannot be shared by two subjects with different total sets of experiences.

Coleman's similar argument[11] involves imagining two subjects, with experiences respectively "pervaded by a unitary blueness" and "pervaded by a unitary… redness" (p. 15), combining into a composite subject:

> To say these points of view were present as components in the experiential perspective of the uber-subject… would therefore be to say that [it] experienced a unitary phenomenal blueness and a unitary phenomenal redness, i.e. had synchronous experiences as of each of these qualities alone, *to the exclusion of all others.* For it is these qualities each *on their own* that characterise, respectively, the perspectives of the original duo. Experience excludes, as well as includes. (Coleman 2013, p. 33, emphasis in original)

I take Coleman's claim that "experience excludes, as well as includes", and that each part's experience is best captured as a certain quality "to the exclusion of all others", as a version of phenomenal holism, emphasising the role of absences. In both arguments, it is alleged that the experiences of the parts *as experienced by the parts* cannot be among the experiences of the whole *as experienced by the whole*,

[11] Coleman does not frame his argument as targeting the sharing of token experiences, and in fact allows that 'qualitative elements' could perhaps be shared between two subjects, perhaps in cases of 'telepathy' (2013, p. 32). What his argument aims to object to is the compounding of 'perspectives', where conscious experience is taken to involve both qualitative elements and a perspective on them; since I want my notion of a token experience to capture both aspects of it, I take it that an argument against the compounding of perspectives is *a fortiori* an argument against the sharing of token experiences, understood as having a perspectival character.

because the whole would have to experience them both with and without the changes in phenomenal character which come from being unified with each other.[12]

My own formulation of the argument is as follows, drawing heavily from Basile's but making explicit a third premise:

Phenomenal Interdependence (PI): The phenomenal character of an experience depends partly on its phenomenal context, i.e. on which experiences it is unified with.

Phenomenal Essentialism (PE): The phenomenal character of an experience is essential to it.

Suppose for *reductio* that we have two distinct subjects, one (S_4) a proper part of the other (S_5), who share a single experience e_4, but who differ in their other experiences—S_5 has experience e_5 but S_4 does not. Suppose moreover that e_4 and e_5 are unified.

1′. Since e_4 and e_5 are unified, PI implies that that S_5, who experiences both, will experience e_4 with a phenomenal character partly dependent on e_5.
2′. Since S_4 only experiences e_4, PI implies that it experiences e_4 with a phenomenal character independent of e_5.
3′. So e_4 will have different phenomenal characters for S_4 and for S_5.
4′. But by PE, this is impossible: a single experience cannot have two phenomenal characters.[13]

The analogous argument directly against between-subjects unity is not too different. It takes the same two principles (PI and PE) as premises, and then proceeds as follows:

Suppose for *reductio* that we have two distinct subjects, S_1 and S_2, who have distinct experiences e_1 and e_2, which are unified with one another.

1″. Since e_1 and e_2 are unified, PI implies that S_1 will experience e_1 with a phenomenal character partly dependent on e_2 (and vice versa for S_2 and e_2).
2″. Since S_1 only experiences e_1 and not e_2, it will experience e_1 with a phenomenal character independent of e_2.
3″. So e_1 will have different phenomenal characters for S_1.
4″. But by PE, this is impossible: a single experience cannot have two phenomenal characters.

I think once the argument is laid out clearly, it becomes obvious that step 2″ is unmotivated—even if we accept both PI and PE, we have no reason to say that S_1 experiences e_1 in a way that fails to reflect e_2. We would only think that if we had

[12] Strictly, this argument does allow for some experience-sharing, but only between subjects who share their entire sets of experiences (as Basile recognises, pp. 110–111). But this kind of experience-sharing would not enable any interesting sort of between-subjects unity.

[13] We might dispense with PE and simply observe that nothing can differ from itself, even in non-essential ways. But without PE it remains open to claim that the experience 'merely appears different' to different subjects, i.e. has multiple phenomenal characters relative to each. PE is what ensures that the experience's *prima facie* self-differing cannot be analysed away.

confused PI, which is a claim about *phenomenal* unity, with a parallel claim about *subject*-unity—that each experience is dependent on which other experiences are had by its subject. Once this confusion is removed, we have no reason not to think that the phenomenal character of S_1′s experiences can depend on S_2′s experiences, experiences which S_1 is not itself having.

This flaw is obvious in the argument against between-subjects unity. But exactly the same flaw, albeit a little harder to spot, is present in the argument against experience-sharing that Basile and Coleman offer. Step 2′, which claimed that the part (S_4) experienced e_4 in a different phenomenal context from the whole (S_5), is completely unmotivated—in particular, it does not follow from PI, which is a claim about phenomenal unity, not subject-unity. Thus the defender of experience-sharing can simply reply that the part and whole have experiences with the same phenomenal context, and hence with the same phenomenal character—it is just that one experiences the context and the other doesn't.

To put this point in the terms of Coleman's argument, the defender of experience-sharing can accept that no composite could have a unified experience of "red-to-the-exclusion-of-(blue-and)-all-else [and] blue-to-the-exclusion-of-(red-and)-all-else" (2013, p. 33). Both experiences would be affected by the other so as to lose their 'to the exclusion of' character. But this can be true also of the parts of the imagined composite: one can experience red, and the other blue, without either experience having the character of 'excluding' the other, precisely because those experiences are unified and thus adjust each other's phenomenal character. Perhaps each does, *when isolated*, experience their redness or blueness as excluding all else, but once they are connected into a unified whole, which experiences the red *as* unified with blueness, and vice versa, so do they. This is simply what Phenomenal Interdependence says.[14]

Should we accept a parallel doctrine to PI, framed in terms of subject-unity? I see no obvious reason to do so. The phenomenological observations that are used to support PI appeal to the relations between experiences themselves, not to their relations to a subject. If disunified consciousness in a single subject is possible, then it seems very implausible that the disunified experiences would automatically alter each other's phenomenal character in the way claimed by PI.

Admittedly, insisting that PI pertains to phenomenal unity, not subject-unity, requires making sense of how one subject's experience could have its phenomenal character altered by another experience which that subject *does not have*. If unified experiences 'interpenetrate' with each other, then a subject who experienced one but not all of a set of unified experiences would somehow be aware of their experiences

[14] Of course, the phrase 'to the exclusion of' may simply report the negative fact that nothing other than red (or blue) is being experienced by that part. But it is unproblematic for this to be true of the part but not the whole. It is a trivial logical point that a whole may have some parts which are not X, but nevertheless be X (a house can have parts which are free of asbestos, yet not be free of asbestos because it has other parts which contain asbestos). And if we read 'to the exclusion of' in this way, we are dealing with trivial logical points, not with phenomenology. Coleman's argument turns on inferring from the negative claim "this component subject is not experiencing any blue" to the phenomenological claim "this component subject is experiencing something as excluding blueness". This is a substantive inference, and the defender of experience-sharing can and should reject it.

are interpenetrating with something they do not experience. Or, in the terms of Coleman's argument, if the whole does not experience either blue or red 'to the exclusion of the other', then each part must somehow experience a single colour that does not seem to them to be 'to the exclusion of' all others.[15] And this may seem hard to make sense of.

But this problem is simply the problem of providing a phenomenologically intelligible account of between-subjects unity. What we are worrying is simply: what could it be like for one of my experiences to be unified with another experience that I do not have? And that worry is the topic of my next section.

5 The argument from weak inconceivability

My third argument against between-subjects unity is not so much a proof of impossibility as a challenge to substantiate its possibility by providing a positive conception. The objector demands that we go beyond simply showing that between-subjects unity is not logically ruled out, towards getting a phenomenological grip on what it might be like for the subjects involved.

Bayne (2010, pp. 43–45) provides a useful way to formulate this argument. Following Van Cleve (1983), Bayne distinguishes 'strong' and 'weak' inconceivability: something is strongly inconceivable to me if it positively strikes me as impossible, but is weakly inconceivable if I merely "cannot see that it is possible" (Bayne 2010, p. 43). The previous two sections considered arguments that between-subjects unity is *strongly* inconceivable, due to implying some contradiction; but there remains the worry that it might be weakly inconceivable, if we can form no positive idea of what it would be for a subject to experience something that was phenomenally unified with something it did not experience. And this might suffice to cast "a significant degree of suspicion" on the notion (Bayne 2010, p. 44—this is what Bayne claims the weak inconceivability of 'partial unity' casts on that notion).

To dispel this cloud of suspicion, I will propose a candidate conception of what between-subjects unity would be like for the subjects involved. In keeping with my method so far, I will accept the subsumption principle and pervasive phenomenal interdependence: the latter, in particular, makes my task significantly harder, since it requires me to accommodate the way that two unified experiences interpenetrate, becoming in some sense 'present in' each other. Yet this must not entail that the subject of one also experiences the other, or we lose the original idea that they are having only one of the two experiences. Somehow each subject must be aware of the other's experience as present in their own, yet not fully given to them. However,

[15] Coleman mentions briefly the possibility that the component perspectives are changed by their combination, but maintains that "this does nothing to get around the problem… if the [component] subjects survive, as they must, there will be three actual experiencings… It's not possible to imagine Ub's unified point of view as comprising two synchronous experiencings…" (2013, p.33, Fn. 38). At this point, however, I am unable to discern the argument, since a supporter of 'perspective-combination' will simply agree that there are three actual experiencings, in the sense of three subjects with three perspectives, and then simply disagree that we cannot regard the third as comprising the other two together. What I have tried to refute is the part of Coleman's discussion which I know how to articulate as an argument.

I believe we have a familiar example of just this kind of phenomenology, in the widespread experience of 'perceptual co-presentation'.

It is a pervasive feature of everyday experience that we perceive objects as having features we do not perceive, so that we are perceptually aware of those features yet in some sense unable to perceive them. One label for this phenomenon is 'amodal perception' (contrasting with 'modal completion', Michotte et al. 1991), reflecting the idea that we somehow perceive co-presented features without sensory stimulation and thus not in any sensory modality. The term 'co-presentation', along with 'adumbration' and 'horizon', are used by authors in the phenomenological tradition (Husserl 1970, 1982a, b, 2001; Merleau-Ponty 1962; cf. Kelly 2004).

The most popular examples involve visual occlusion, as when seeing three-dimensional objects with front and back side. When I look at a coffee cup, there is a sense in which I see only the front of it, but also a broader sense in which I see the cup itself, a whole with a front and back. I am in some indirect sense aware of the back of the cup, but simultaneously aware that I am not aware of it in the same sense that I am aware of the front; it is co-presented but not presented. Moreover, I am aware of the back *in virtue of* being aware of the front, which I see *as* the front *of* something which also has a back. Let us say that the front is, for me, the 'revealed aspect', and the back the 'concealed aspect'; while I am aware of the former as 'given', I am aware of the latter only in a weaker sense, as 'not given'.

Consider also the table which the cup stands on: I perceive its surface as wide and unbroken even though part of that surface is blocked out the cup. Some sections are presented to me through vision and they co-present the other sections as more sections of the same continuous surface. Or consider cases where poor visibility stops us from seeing something clearly. Rather than experiencing the obscured object as somehow fuzzy in itself, we experience it as having plenty of detail, which *we* cannot make out. Insofar as we are aware of this detail as not visible, it is co-presented to us. Here what is revealed and concealed are not spatially separate parts, but the broad outlines and fine details of a single object. The rough aspect co-presents the detailed aspect, by being experienced *as* a rough and imperfect view of an object that can be seen in better ways.

One important dimension along which cases of perceptual co-presentation vary is their informativeness. What is directly perceived tells us about what is not given but co-presented. I may not be able to see the rear side of my coffee cup, but I can tell a fair bit about its overall size, shape, and location, just from the aspect which is revealed. I can also estimate its likely colour, patterning, etc.[16] Different cases may give more or less information, or may give more information to subjects with more background knowledge. For example, a brief glimpse through a window at night may reveal a slight movement, experienced as that of some vast creature without yielding any definite information as to its nature or shape—by contrast, when the same vast creature is seen in daylight, it will still have concealed aspects (e.g. its

[16] Thus perceptual co-presentation is in some ways a form of 'deferred perception' (Dretske 1994), but not all cases of deferred perception are perceptual co-presentation—when paleontologist 'sees that the hadrosaur was attacked by a tyrannosaur', just by looking at the bones, they have no sense of the tyrannosaur as perceptually present to them.

opposite side) but their nature will be much more closely specified by what is then revealed.

Philosophers and psychologists have done a lot of work on perceptual co-presentation (e.g. Michotte et al. 1991; Clarke 1965; Matey 2013), addressing various questions about its phenomenology (e.g. Merleau-Ponty 1962; Kelly 2004), about its neural basis and relationship to certain visual illusions (Kellman and Shipley 1991; Ramachandran et al. 1995), or about which representational faculty it involves—perception, cognition, imagination, or something else (Noë 2005; Nanay 2010; Briscoe 2011).[17] I wish to be neutral on all these points. What matters is just that in perceptual co-presentation we are aware of something specifically as not fully given to us, via its connection to something which is given to us. This makes it perfectly suited to explaining the phenomenology of between-subjects unity.

So what is it like to be one end of a between-subjects unity relation? I suggest that it is best thought of as a sort of 'co-presentation' of another subject's experience, by one's own experience. To put it another way, it would involve experiencing one's own experience *as* the revealed aspect of something of which other experiences were concealed aspects.[18]

This proposal faces three immediate challenges. The first concerns whether we are really *aware of* our experiences, for otherwise we cannot be aware of them *as* the revealed aspects of something. This is a deeply perplexing question, on which diametrically opposed positions have been regarded as introspectively evident (see, e.g., Block 2003; Kriegel 2009; Strawson 2013). But I think for our purposes here we could finesse the issue by simply claiming the following parallel: we are aware of our experiences in whatever sense we are aware of the unity among our experiences.

In particular, according to the doctrine of the 'transparency' of experience, we are not, properly speaking, 'aware of' our experiences but only of the contents they present to us, and can ascribe experiences to ourselves introspectively only by a form of 'deferred perception' whereby we learn *that* we are having a certain experience by being aware of its content and employing theory-of-mind concepts (Dretske 1994). The unity of consciousness would then have likely be accounted for somehow within the contents of experience, e.g. by taking unified experiences to have conjunctive content, or to immediately represent their objects as standing in some relation (cf. Masrour 2014, also Bayne and Chalmers 2003, p. 53ff, Tye 2003, p. 38ff). The proposal under consideration would then have to be formulated in a corresponding way. Perhaps between-subjects unity is experienced as the content of my experience co-presenting the content of your experience, but differently from

[17] This latter question has been connected to the broader idea, particularly associated with Strawson's (1974) defence of Kant, that all perception is in some sense infused with imagination.

[18] Note that this stops short of giving a full analysis of what the unity relation in such cases would *require*: saying how it would be experienced by each subject when it does obtain is not giving necessary and sufficient conditions for its obtaining. The latter might well include strictly relational facts, such as causal, informational, or referential links between how the relation is experienced by one subject and how it is experienced by the other, such as the requirement that each subject be co-presented with the other's experiences as a result of the other having those experiences, or of the other being likewise aware of the first's.

the normal case where one content indicates another because the fact that what is co-presented is co-presented *as* the content of another experience—which brings us neatly to the second challenge to my proposal.

The second challenge is whether experiences can be co-presented, i.e. whether we can be aware of them as concealed. To support a positive answer, consider our awareness of other people's experiences through their outward behaviour, which I think is better analysed as a case of perceptual co-presentation than as either a case of reasoning or a case of direct perception.[19]

For instance, when two people conversing see each other's facial expressions, or hear each other's voices, they may well perceive these expressive actions as the revealed aspect of a conscious state which is not directly given, but is nevertheless experienced as 'present' through its expression. We do not generally feel as though we have to infer someone's anger from their words and expressions, in any highly intellectual way. Rather, we 'see the anger in their face'—and yet we also tend to think that the anger itself is private: we are aware that it is fully given only through this other experience that we do not have.

Here an external perceptual object—a face, a voice, a gesture—co-presents an experience. The phenomenology I am proposing for between-subjects unity is importantly different: an experience is co-presented not by an external perception but through our awareness of one of our own experiences. It would involve my experiences giving me insight into how the world seems to someone else, in a way that came to me as automatically and immediately as my own experiences do: it would be, we might say, a sort of 'introspective co-presentation'.[20] But we should be clear that this is not a matter of each subject *undergoing* the other's experience is a special way: the whole point of between-subjects unity is that the subjects do not share all of their experiences. Rather, each subject undergoes *their own* experiences in a way that provides them with an indirect awareness of the other's.

Perhaps the best model we have for this sort of connection would involve people with whose viewpoints we have become so intimately familiar, and so constantly concerned, that we cannot think of or perceive anything without it occurring to us how they would think of or perceive that thing. In cases like this—which typically involve parents, teachers, or mentors of some other kind—we might naturally say that we have the other person 'in our heads'. But this kind of relation to another person's thoughts is not in virtue of any actual relation to their present experiences,

[19] This approach to mind-perception is defended by Smith (2010, 2015), drawing on ideas in Husserl (1982a, b). I defend a slightly different version of the same approach in other work.

[20] Here again the proposal would have to take slightly different forms to accommodate different views on what it means to be aware of an experience. If experience is transparent, and my awareness of it is by the application of theory-of-mind concepts to experiential contents, then presumably perception of others as having experiences would involve applying such concepts to the observed behaviour of others—a different mode of application from the introspective one. My proposal for between-subjects unity would then say that it involves learning of another subject's experiences through applying theory-of-mind concepts, in the manner typical of self-knowledge rather than other-knowledge, to co-presented contents. Each subject would, somehow, be co-presented with certain contents in such a way as to put them in a position, given the right conceptual abilities, to thereby know that someone is having an experience which directly presents that content to them.

but in virtue of a sort of 'internalisation', whereby we have constructed a working simulation of them in our own heads. (It could, after all, persist even after they themselves have died). For our present experiences to be really unified with theirs, we would have to have the same sort of automatic and immediate awareness of how things seem to them, but have it in virtue of a present flow of information.

The third challenge arises as follows. Between-subjects unity and within-subject unity are meant to be the same relation, and so whatever we say about the experience of the former will also force us to say something similar about the latter. Yet it seems contradictory to think of within-subject unity in terms of experiences being co-presented. How can my awareness of one experience involve 'awareness of not being aware' of another, if in fact I *am* aware of both? But we must be careful how we characterise the negation in 'awareness of not being aware'. We might gloss it as awareness that:

- "This is not given at all, to anyone."
- "This is not given to me."
- "This is not given to me through this particular experience."

The third of these is preferable, because something can be both presented to us and co-presented at the same time, with no sense of contradiction or impossibility. I may see a three-dimensional object as having a concealed rear side, while also seeing that rear side directly in a mirror placed behind it. If I am right that the negative aspect of perceptual co-presentation is best captured as "this aspect is not given through *this* experience", then a subject whose experiences mutually indicate each other in this way need not think that they both are and are not being given each aspect, but merely that each aspect is given in one experience, and intimated but not given in the others.[21]

6 Conclusions

I began this paper with an under-explored question: could the unity we find in our own experience obtain across the boundaries between subjects: could one subject's experience stand in that relation to another's? I have argued, firstly, that we cannot show this to be impossible, unless we endorse the exclusivity principle in its strong form, along with a subsumptive account of unity; I then argued that we have no good reason for preferring the strong form of exclusivity that rules out between-subjects unity to the weak form that allows it. I argue, moreover, that between-subjects unity is compatible with interdependence among the phenomenal characters of unified experiences, and that we may have a way to make phenomenological sense of it through the familiar experience of perceptual co-presentation: between-subjects unity, we might suppose, would involve a sort of "introspective co-presentation".

[21] One might even take this sort of mutual-co-presentation as an *analysis* of unity, either a rival or a supplement to others like the subsumption analysis. But to fully lay out and defend such a view would require a paper of its own.

Obviously there is room for much more to be said on each of these three points: finding the right account of phenomenal unity, determining the possibility of experience-sharing, and fleshing out the idea of introspective co-presentation. But I hope the arguments and proposals of this paper indicate a fertile space for future investigation, especially by philosophers whose independent interests—such as in panpsychism, or phenomenal holism, or the potential for collective mentality—give them reason to be interested in whether the unity of consciousness can exist not only within a subject, but between subjects.

Compliance with ethical standards

Ethical statement This work was produced in conformity to all the ethical standards outlined here http://www.springer.com/philosophy/journal/11098. No research was performed on human or animal subjects. The only potential conflict of interests is a $500 essay prize awarded to a fore-runner of this paper by the Canadian Philosophical Association at their 2014 annual meeting.

References

Avicenna. (1952). Avicenna's psychology. In *The book of healing* (F. Rahman, Trans.). Westpoint, CT: Hyperion Press. (**Originally published 1027**).

Baker, L. R. (1999). Unity without identity: A new look at material constitution. *Midwest Studies in Philosophy, 23*, 144–165.

Barnett, D. (2008). The simplicity intuition and its hidden influence on philosophy of mind. *Nous, 42*(2), 308–335.

Basile, P. (2010). It must be true—But how can it be? Some remarks on panpsychism and mental composition. *Royal Institute of Philosophy Supplement, 67*, 93–112.

Bayle, P. (1991). *The historical and critical dictionary: Selections*. Translated and introduced by Richard Popkin, with assistance from Craig Brush. Indianapolis: Hackett. (Original Publication 1697).

Bayne, T. (2010). *The unity of consciousness*. Oxford: Oxford University Press.

Bayne, T., & Chalmers, D. (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, dissociation* (pp. 23–58). Oxford: Oxford University Press.

Block, N. (1978). Troubles with functionalism. In C. W. Savage (Ed.), *Perception and cognition: Issues in the foundations of psychology* (pp. 261–325). Minneapolis: University of Minnesota Press.

Block, N. (2003). Mental paint. In M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler Burge*. Cambridge: MIT Press.

Brentano, F. (1987). *The existence of god: Lectures given at the Universities of Würzburg and Vienna, 1868–1891* (S. Krantz, Ed., Trans.). Dordrecht: Martinus Nijhoff Publishers. (**Originally delivered 1868–1891**).

Briscoe, R. (2011). Mental imagery and the varieties of amodal perception. *Pacific Philosophical Quarterly, V92*(2), 153–173.

Butler, J. (1860). *The analogy of religion, natural and revealed, to the constitution and course of nature*. Edited, with an analysis, by J. T. Champlin. Boston: J. P. Jowett and Co. (**Originally published 1736**).

Chalmers, D. (2015). Panpsychism and panprotopsychism. In T. Alter & Y. Nagasawasa (Eds.), *Consciousness in the physical world: Essays on Russellian Monism*. Oxford: Oxford University Press.

Chalmers, D. (Forthcoming). The combination problem for panpsychism. In G. Brüntrup & L. Jaskolla (Eds.), *Panpsychism*. Oxford University Press.

Chudnoff, E. (2013). Gurwitsch's phenomenal holism. *Phenomenology and the Cognitive Sciences, 12*(3), 559–578.

Clarke, T. (1965). Seeing surfaces and physical objects. In Max Black (Ed.), *Philosophy in America*. Ithaca, NY: Cornell University Press.

Clarke, S. (1978). *The works*. New York: Garland Publishing. **(originally published 1738)**.

Coleman, S. (2012). Mental chemistry: Combination for panpsychists. *Dialectica, V66*(1), 137–166.
Coleman, S. (2013). The real combination problem: Panpsychism, micro-subjects, and emergence. *Erkenntnis, 79*(1), 19–44.
Dainton, B. (2000). *Stream of consciousness: Unity and continuity in conscious experience*. London: Routledge.
Dainton, B. (2008). *The phenomenal self*. Oxford: Oxford University Press.
Dainton, B. (2010). Phenomenal holism. *Royal Institute of Philosophy Supplement, 67*, 113–139.
Dainton, B. (2011). Review of consciousness and its place in nature. *Philosophy and Phenomenological Research, V83*(1), 238–261.
Descartes, R. (1985). Meditations on first philosophy. In *The philosophical writings of descartes* (Vol. 2) (J. Cottingham, R. Stoothoff & D. Murdoch, Eds., Trans.). Cambridge: Cambridge University Press. (**Originally published 1641**).
Dretske, F. (1994). Introspection. *Proceedings of the Aristotelian Society, 94*, 263–278.
Goff, P. (2006). Experiences don't sum. *Journal of Consciousness Studies, 13*(10–11), 53–61.
Gurwitsch, A. (1964). *The field of consciousness*. Pittsburgh: Duquesne University Press.
Huebner, B. (2011). Genuinely collective emotions. *European Journal for Philosophy of Science, V1*(1), 89–118.
Husserl, E. (1970). *Logical investigations* (J. N. Findlay, Trans.). London: Routledge and Kegan Paul.
Husserl, E. (1982a). *Ideas pertaining to a pure phenomenology and to a phenomenological philosophy, first book* (F. Kersten, Trans.). Dordrecht: Kluwer Academic Publishers.
Husserl, E. (1982b). *Ideas pertaining to a pure phenomenology and to a phenomenological philosophy: First book* (F. Kersten, Trans.). The Hague: Nijhoff. (**Originally published in German 1913**).
Husserl, E. (2001). *Analyses concerning active and passive synthesis: Lectures on transcendental logic*. Lectures given between 1920 and 1926 (A. Steinbock, Trans.). Springer.
James, W. (1890). *The principles of psychology*. New York, NY: Harvard University Press.
Kant, I. (1998). *Critique of pure reason* (P. Guyer & A. Wood, Trans.). Cambridge: Cambridge University Press. (**Originally Published 1781**).
Kellman, P., & Shipley, T. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology, 23*, 141–221.
Kelly, S. (2004). On seeing things in Merleau-Ponty. In T. Carmon (Ed.), *The Cambridge companion to Merleau-Ponty*. Cambridge: Cambridge University Press.
Knobe, J., & Prinz, J. (2008). Intuitions about consciousness: Experimental studies. *Phenomenology and Cognitive Science, 7*, 67–83.
Koksvik, O. (2014). Three models of phenomenal unity. *Journal of Consciousness Studies, 21*(7–8), 105–131.
Kriegel, U. (2009). *Subjective consciousness: A self-representational theory*. Oxford: Oxford University Press.
Lennon, T., & Stainton, J. (Eds.). (2008). *The Achilles of rationalist psychology. Studies in the history of philosophy of mind* (Vol. 7). Berlin: Springer.
Lewis, D. (1976). Survival and identity. In A. Rorty (Ed.), *The identities of persons*. Berkeley: California.
Lotze, H. (1894). *Microcosmus: An essay concerning man and his relation to the world* (4th ed., Vol. I) (E. Hamilton & E. E. Constance Jones, Trans.). New York: Charles Scribner's Sons. (**Originally published in German 1856**).
Masrour, F. (2014). Unity of consciousness: Advertisement fora Leibnizian view. In C. Hill & D. Bennett (Eds.), *Sensory integration and the unity of consciousness* (pp. 323–346). Cambridge, MA: MIT Press.
Matey, J. (2013). Representing the impossible. *Philosophical Psychology, V26*(2), 188–206.
Mendelssohn, M. (1767/1789/2002). *Phaedon or, The Death of Socrates.* In (C. Cullen, Trans.), Schiller Institute Inc. Accessed July 2013 http://www.schillerinstitute.org/transl/mend_phadn_cullen.html (**Originally published 1767**).
Merleau-Ponty, M. (1962). *The phenomenology of perception* (C. Smith, Trans.). London: Routledge and Kegan Paul. (**Originally published 1945**).
Merricks, T. (2001). *Objects and persons*. Oxford: Oxford University Press.
Michotte, A., Thinès, G., & Crabbé, G. (1991). Amodal completion of perceptual structures. In G. Thinès, A. Costall, & G. Butterworth (Eds.), *Michotte's experimental phenomenology of perception* (pp. 140–167). Hillsdale, NJ: Erlbaum. (**Originally published in French 1964**).

Mijuskovic, B. L. (1984). *The Achilles of rationalist arguments: The simplicity, unity, and identity of thought and soul from the Cambridge Platonists to Kant: A study in the history of an argument*. The Hague: Martinus Nijhoff.
Nagel, T. (1971). Brain bisection and the unity of consciousness. *Synthese, 22*, 396–413.
Nanay, B. (2010). Perception and imagination: Amodal perception as mental imagery. *Philosophical Studies, 50*, 239–254.
Noë, A. (2005). Real presence. *Philosophical Topics, 33*, 235–264.
Parfit, D. (1971). Personal identity. *Philosophical Review, 80*, 3–27.
Parfit, D. (1999). Experiences, subjects, and conceptual schemes. *Philosophical Topics, 26*, 217–270.
Plotinus. (1956). *Enneads.* In (S. Mackenna, B. S. Page, Eds., Trans.). Faber and Faber limited. http://archive.org/stream/plotinustheennea033190mbp#page/n11/mode/2up (accessed july 2013).
Proclus. (1963). *The elements of theology.* In (E. R. Dodds, Ed., Trans.). Oxford: Clarendon Press.
Putnam, H. (1965/2003). The nature of mental states. In T. O'Connor, D. Robb, J. Heil (Eds.), *Philosophy of mind: Contemporary readings*, pp. 223–235. Routledge.
Ramachandran, V. S., et al. (1995). Filling in the gaps in logic: Reply to Durgin et al. *Perception, 24*, 841–843.
Roelofs, L. (2014). Phenomenal blending and the palette problem. *Thought, 3*, 59–70.
Roelofs, L. (2015). *Combining minds: A defence of the possibility of experiential combination.* Ph.D. Thesis, University of Toronto.
Schachter, J. (2002). Pierre Bayle, matter, and the unity of consciousness. *Canadian Journal of Philosophy, V32*(2), 241–265.
Schechter, E. (2010). Two unities of consciousness. *European Journal of Philosophy, V21*(2), 1–25.
Schwitzgebel, E. (2014). If materialism is true, the United States is probably conscious. *Philosophical Studies, 172*, 1697–1721.
Seager, W. (1995). Consciousness, information and panpsychism. *Journal of Consciousness Studies, V2–3*, 272–288.
Searle, J. (2000). Consciousness. *Annual Review of Neuroscience, 23*, 557–578.
Shoemaker, S. (2003). Consciousness and co-consciousness. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation*. Oxford: Oxford University Press.
Smith, J. (2010). Seeing other people. *Philosophy and Phenomenological Research, 81*(3), 731–748.
Smith, J. (2015). The phenomenology of face-to-face mindreading. *Philosophy and Phenomenological Research, 90*(2), 274–293.
Sprigge, T. (1983). *The vindication of absolute idealism.* Edinburgh: Edinburgh University Press.
Strawson, P. F. (1974). Imagination and perception. In P. F. Strawson (Ed.), *Freedom and resentment and other essays* (pp. 45–65). London: Methuen.
Strawson, G. (2003). What is the relation between an experience, the subject of the experience and the content of the experience? *Philosophical Issues, 13*(1), 279–315.
Strawson, G. (2006). Realistic monism: Why physicalism entails panpsychism. In A. Freeman (Ed.), *Consciousness and its place in nature: Does physicalism entail panpsychism?*. Exeter: Imprint Academic.
Strawson, G. (2013). Self-intimation. *Phenomenology and the Cognitive Sciences, 14*, 1–31.
Tononi, G. (2012). Integrated information theory of consciousness: An updated account. *Archives Italiennes de Biologie, 150*(2–3), 56–90.
Tye, M. (2003). *Consciousness and persons: Unity and identity*. Cambridge, MA: MIT Press.
Unger, P. (1990). *Identity, consciousness, and value*. Oxford: Oxford University Press.
Unger, P. (2006). *All the power in the world.* Oxford: Oxford University Press.
Van Cleve, J. (1983). Conceivability and the Cartesian argument for dualism. *Pacific Philosophical Quarterly, 64*, 35–45.
Zuboff, A. (1990). One self: The logic of experience. *Inquiry, V33*(1), 39–68.