

In Defence of Phenomenal Sharing

Philip Goff and Luke Roelofs

Section 1: Sharing and Exclusivity

William James expresses a commonplace idea when he writes that:

Each of [our] minds keeps its own thoughts to itself.... No thought even comes into direct sight of a thought in another personal consciousness than its own. Absolute insulation... is the law... The breaches between such thoughts are the most absolute breaches in nature. (James 1890, p.226)

This idea - that each individual consciousness is absolutely insulated from all others - could be unpacked in a variety of ways, but the strand we wish to focus on here is the denial of what we will call 'phenomenal sharing'. James says that each mind 'keeps its own thoughts to itself': the opposite possibility, that he denies, would presumably be for one thought to be shared between two minds. As we will discuss, there are a number of contexts where philosophers might feel driven to postulate such a 'sharing' of mental particulars, but there are a number of arguments and objections that seem to show it impossible. We believe that none of these arguments are decisive: while we cannot positively establish the possibility or actuality of mental sharing, we hope to show that philosophers who have independent reasons to postulate it in particular cases need not hold back from doing so.

Because many of the arguments against shared mental particulars turn specifically on facts about phenomenal consciousness, we will focus our attention on phenomenal sharing - what might be called the sharing of 'experiences' (though see section 2 for some terminological complications). Our target could thus be put as a claim of 'Exclusivity': subjects have exclusive ownership of their experiences. Dainton endorses an idea like this, writing that "it seems plausible to think that subjects and experiences are governed by an Exclusivity Principle along these lines:"

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 . (Dainton 2011, p.246)

Our disagreement with this Exclusivity Principle is not total: we agree that there is a sense in which the experiences of one subject cannot simultaneously belong to ‘another’ subject. But there are two ways to understand what is ‘another’ subject: we might mean simply a ‘distinct’, i.e. nonidentical, subject (as in Dainton’s formulation), or we might mean ‘discrete’, i.e. nonoverlapping, subjects, two subjects that do not share any parts. What it means for two subjects to share parts depends somewhat on how subjects are conceived, in ways that we explore further in subsection 3.3. But for now we can think of overlap in physical/spatial terms: a subject is either a physical system that is conscious, or something arising from a physical system, and two subjects overlap if their physical systems overlap, i.e. share physical parts.

Taking these two readings of the Exclusivity Principle, together with their negations, 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.

Weak Sharing is the negation of Strong Exclusivity, while Strong Sharing is the negation of Weak Exclusivity; Strong versions of both Exclusivity and Sharing entail weaker versions. But, crucially, Weak Exclusivity and Weak Sharing are compatible – experiences might be shareable by non-discrete, i.e. overlapping, subjects, but not shareable between discrete subjects. Indeed, this is the position we will defend.

There are certain **major** questions here that we are leaving open. How spatially ‘big’ are experiences - that is, how extensive is the neural basis for a given experience? Indeed, can the complexity of our phenomenal field be divided into component experiences at all? And how much overlap between two subjects is compatible with counting them as still *two*, as opposed to being the very same subject? How someone answers these questions will affect the range of cases in which phenomenal sharing is even a possibility. **At one extreme, some** of the applications we discuss in section 4 are for views like panpsychism, on which simple experiences have *extremely* small physical bases, or for too-many-minds problems, which usually assume quite fine-grained individuation of subjects. Both of these starting points imply wide scope for experience sharing. **But some defenders of experience sharing consider arrangements with much more extensive overlap: Hirstein (2008, 2012), for example, focuses on the idea of duplicating the brain’s internal connections so as to connect two people’s frontal cortices to a single set of sensory cortices, so that two minds share their entire sensory consciousness. And in the extreme,** if we thought both that any given human experience requires some large portion of the brain to occur, and also thought that two human brains which overlapped enough to share that large portion would *ipso facto* constitute a single subject, not two, then phenomenal sharing might not make sense in any realistic human cases. But we feel the question of its in-principle possibility remains an interesting one, on **any of these positions.**¹

Part of our reason for defending only the ‘Weak’ sort of phenomenal sharing is that something analogous seems to be the right view of physical properties: they seem to be shareable between overlapping physical objects, but not between discrete ones. For example, although a cup weighs 100g, and its handle weighs 20g, there is not 120g of mass here: the weight of the handle already belongs to the whole cup, and so should not be added to it. More broadly, the cup seems to share the spatial location(s), colour(s), powers, etc. as its other parts; if there is a red spot, or a scratch, or some other feature, on one

¹ Thanks to Ned Block for pushing us on this point.

side, both the side and the whole cup instantiate ‘red-spotted-ness’ (or ‘scratched-ness’, etc.), and yet ‘red-spotted-ness’ is only instantiated once, not twice over.

On the other hand, it does not in general seem that two discrete substances can simultaneously share a property instance. If we try to imagine, for instance, a single instance of redness belonging to two discrete surfaces, it seems that all we can imagine is a token of redness, some parts of which belong to the one surface, and other parts of which belong to the other.

So our aim is to defend the principle of Weak Sharing: to the extent that two subjects overlap - one containing the other as a proper part, or both sharing a single proper part - they may share particular experiences. In section 2, we will address certain definitional challenges to this idea, that will require us to be careful about how we use the term ‘experience’. In section 3, we consider 5 objections to the possibility of phenomenal sharing, showing that all can effectively be resisted, though doing so often requires accepting substantive assumptions about the nature of experience, subjects, and phenomenal unity.

This conclusion would not, of course, have immediate consequences for our everyday picture of human beings, on which each is a single subject, fully discrete from all the others. But, as we explain in section 4, opening up this possibility may have far-reaching implications for a number of debates in the philosophy of mind. Some of these concern how to think about hypothetical or unusual cases where two brains are physically connected; some concern how to think about the relationship between overlapping physical entities (organisms, heads, brains, and so on) in ordinary cases; some concern theories of consciousness on which parts of the human brain may be conscious in their own right.

Section 2: Experiences, Phenomenal Tokens, and Phenomenal Havings

There is a risk of this debate devolving into an argument over terminology, or being closed off prematurely by being defined out of existence. For on the one hand, everyone agrees that it is possible to

share phenomenal properties in the *type* sense - that is, two subjects may have experiences of exactly the same type. And on the other hand, any phenomenal particular that is individuated by the unique subject that bears it will be unshareable by definition. So if, for example, we adopted the tripartite account of experience individuation proposed by Bayne (2010, pp. 24-29), on which experiences are individuated by their phenomenal character, subject, and time, then it will be contradictory for one 'experience' to belong to two distinct subjects. Indeed, this is one plausible way to individuate property instances generally - by the property instantiated, the object that instantiates it, and the time of instantiation (see, e.g. Ehring 1996, 462). If experiences are individuated like this, then the experiences of any two subjects must be accounted different experiences just in virtue of their different subjects. So it might look at first as though phenomenal sharing is a non-starter: either it is trivial, if it means sharing of universal types, or it is contradictory, if it means sharing of subject-individuated tokens.

But this sort of objection has nothing specifically to do with experience: all the same things could be said about physical properties. Above we noted that intuitively, overlapping physical objects do not duplicate their properties, but rather share them - and consequently they do not compete for causal efficacy nor add up in calculations of parsimony. But if we individuate property instances by the property instantiated, object instantiating it, and time of instantiation, then this intuitive idea will become contradictory. Even if we accepted this, it would be a rather underwhelming reason to reject phenomenal sharing. The idea we aim to challenge is that there is something special about phenomenal properties that makes them uniquely unshareable. So we should distinguish three different sorts of entity:

1. Property types
2. Property tokens, individuated in a way that does not rule out their being shared
3. Property tokens, individuated in a way that rules out their being shared

In a physical case, for instance, these three might be, respectively: 'the property of being scratched', in general; a particular instance of this property, say that on the side of my cup, which we might simply call

'the scratch'; and things like 'the fact of the cup being scratched', 'the fact of the left side being scratched', 'the fact of this region of the left side being scratched', and so on. The first of these is a single thing present wherever anything is scratched; the second is a particular at a specific time and place, but still intuitively one in number; the third are numerous - there is one for each of the overlapping objects that have the scratch. We interpret Exclusivity as saying that when it comes to phenomenal properties, even entities of type-2 are not shareable - that phenomenal properties cannot be shared *in the way that physical properties can*. And we aim to challenge this thesis: we argue that phenomenal properties can be shared in the way that physical properties can, because individual phenomenal entities of type 2 can belong to more than one subject.²

There is room for different views on both the metaphysics and the semantics of these three types of entity. We might try to construe them as tropes, as instantiations of universals, or as events, depending on how we think about their relationship to time. We might think that type-2 entities are more basic than type-3 entities: reflecting on physical cases like my cup might make it natural to say that definitionally-unshareable things like the 'the left side having the scratch' are just abstractions from the more basic reality of the single scratch, or even just artefacts of our language. Alternatively, we might think that type-3 entities are more basic, and that type-2 entities are best analysed as sets of type-3 entities related in certain ways, e.g. the scratch is a set of connected havings-of-scratchedness. That is, we might say that while *in metaphysical strictness*, there are different property instances for each scratched object, they are related in such a way that it makes sense for us to speak of them as one. There is some sort of special metaphysical intimacy between them such that, given one instance, the instances possessed by its overlappers are nothing over and above that one. For example, Sutton (2014) analyses this intimacy as

² It is worth stressing that we are interested in the states of the subject, not in any objects that the subject is aware of or otherwise related to, whether those objects are intra-mental or extra-mental. So our claim is stronger than that of, e.g., Coleman 2014, who analyses experiences as awareness of qualities, and allows that two subjects might be aware of the same token quality while denying that they can thereby be having the same experience of it. Our thanks to Sam Coleman and Ali Rezaei for pushing us on this point.

“shar[ing] a supervenience base for” the relevant properties (p. p.622): if the part of something whose intrinsic features are sufficient for a given property to be instantiated is shared by another whole, then both wholes will instantiate that property in a “non-summative” way, meaning although two beings instantiate the property, the total amount of that property is not thereby increased (p.622). Sutton proposes this analysis in part to allow for “two non-identical beings that think the same thoughts” (p.622), as a solution to the ‘Too Many Minds’ problems we briefly discuss in subsection 4.2. Though Sutton’s supervenience-based analysis is congenial to us, in this paper we aim for neutrality about the metaphysics of properties: our aim is simply to evaluate arguments that, however we might decide to analyse the intuitive cases of physical sharing, something stops us extending that analysis to allow for phenomenal sharing.³

There is also room for different ways of labelling phenomenal entities of the different types. In particular, the term ‘experiences’ could be used to describe either type-2 or type-3 phenomenal particulars. Indeed, the present authors differ on this point: Roelofs prefers to call type-2 entities ‘experiences’, and thus say that experiences can be literally shared, while Goff prefers to reserve the term ‘experiences’ for type-3 entities, and thus describes our thesis not strictly as experience-sharing but as phenomenal-property-sharing. (This view might still be called ‘experience-sharing’ in an indirect sense: we might simply stipulate that two subjects ‘share an experience’ when their experiences are based in the same phenomenal property instance.) To forestall confusion over words, in what follows we will avoid the word ‘experiences’ in favour of more neutral descriptions of the entities in question:

Type 1: Phenomenal types (or ‘phenomenal property types’)

Type 2: Phenomenal tokens (or ‘phenomenal property tokens’)

³ Some philosophers might feel nominalist scruples against any attempt to reify properties this way; it is then up to them to offer an account of the sense in which the cup being scratched is nothing over and above its side being scratched. This might be done, for instance, in terms of grounding or entailment relations among facts: the fact that the cup is scratched follows in a certain way from the side being scratched, or is grounded in the same underlying facts as the fact that the side is scratched, or both facts symmetrically ground the other, or something else. Whatever analysis is favoured, we maintain that there is no special objection to extending it to phenomenal matters. Our thanks to Davids Builes and Chalmers for pushing us on this point.

Type 3: Phenomenal havings (or ‘havings of phenomenal properties by a subject’)

In this vocabulary, we aim to defend the idea that phenomenal tokens can belong to more than one subject - that they can be shared, just as physical property tokens can be. We may put this as follows:

Phenomenal Sharing: One and the same phenomenal token may belong simultaneously to two or more subjects, if they mereologically overlap.

Section 3: Five Arguments Against Phenomenal Sharing

Let us first consider whether there are any direct arguments for the impossibility of phenomenal sharing.

We will consider 5, which we will term the privacy argument, the intrinsicness argument, the identity argument, the **unity** argument, and the incompatible contexts argument.

Subsection 3.1: The Privacy Argument

The first objection we will consider concerns the seemingly self-evident ‘privacy’ of phenomenal tokens. It seems to be a truism about consciousness that each of us knows our own directly, and others’ only indirectly. This epistemic asymmetry between the way we know our own minds and the way we know other people’s seems almost foundational for the very idea of a conscious mind, and so it would be a major strike against phenomenal sharing if that thesis conflicted with it. And it seems to so conflict: if each element of my experience is ‘private,’ knowable directly only from one perspective, how can that same element also belong to, and thus be knowable from, another perspective? And if we allow them to be shared, and thus knowable to many, won’t we lose sight of what distinguishes ‘subjective’ states from the ‘objective’, ‘public’, facts about material things? While it may be debated how exactly to spell out this

difference between the private and the public, we might be wary of views which remove the contrast altogether.

Hirstein (2008, 2012) actually embraces this consequence: the possibility of technological ‘mindmelding’, as he terms it, refutes the supposed privacy of the mental, and thereby removes a key intuitive objection to physicalism. Conscious experiences, he concludes, are publicly knowable, just like any physical property.

However, Phenomenal Sharing need not remove the intuitive contrast between the subjectivity of experience and the objectivity of non-experiential facts. Consider the following rough statement of the privacy of experience:

Privacy: A phenomenal token is directly knowable by one subject, and knowable only indirectly by any other subject.

Just as we above distinguished Strong Exclusivity from Weak Exclusivity by reading ‘other’ either in terms of discreteness (non-overlap) or in terms of distinctness (non-identity), we can also distinguish Strong Privacy from Weak Privacy (cf. Roelofs 2019, p.63 ff):

Strong Privacy: A phenomenal token is directly knowable by one subject, and knowable only indirectly by any distinct subject.

Weak Privacy: A phenomenal token is directly knowable by one subject, and knowable only indirectly by any discrete subject.

In line with our rejecting Strong Exclusivity but accepting Weak Exclusivity, we likewise reject Strong Privacy but accept Weak Privacy. Doing so allows for experience sharing among overlapping subjects, but it does not remove the original contrast we recognised between the private and the public. Given weak

privacy, I can have direct access to someone's experiences if and only if I either am them or overlap with them (sharing a part, or being part of them, or containing them as a part); thus my identity is crucial to my ability to know their experiences. Nothing comparable is true of their physical properties, which are entirely public, theoretically available equally to any subject in the universe. To know somebody's size or shape, for instance, just requires gaining certain evidence, and in principle anyone can do that, whether or not they are a part of that person, or contain them as a part. A single instance of red-spottedness might be shared between my cup and its handle, but this red-spottedness is not therefore knowable only to them.

We suggest that fundamentally, what explains the privacy of experience is that we know our conscious states by having them. This makes the epistemology of consciousness profoundly different from that of other things, and continues to do so even if we suppose that some phenomenal tokens might be had by more than one subject at once.

Subsection 3.2: The Intrinsicness Argument

A second objection (drawn from correspondence with Sam Coleman) focuses on the idea that 'experiences' are intrinsically conscious, and that being conscious requires being had by a subject: putting these two ideas together, it might seem that a given phenomenal token, if it is to be really something *phenomenal*, must be intrinsically something had by a subject, and which is therefore essentially tied to that subject. And it might seem to follow that anything which is shared between two subjects cannot be essentially tied to either of them, and so cannot be really phenomenally conscious in itself.

This argument might be sidestepped simply by denying the starting point, the idea that consciousness is an intrinsic property of certain events (this is, in fact, Coleman's own position). For example, on a higher-order-thought theory of consciousness, a conscious state is simply one which we are

conscious of (Rosenthal 1986, 2005), i.e. one which is the target of an external state of higher-order awareness. On views like this, there is a profound sense in which ‘what we experience’ is not intrinsically conscious, but just something which we make conscious by becoming aware of it. Neither of the present authors are attracted to this way out: our interest is in whether phenomenal sharing is possible, even given robust realism about consciousness, which usually goes with the thought that conscious events are intrinsically conscious.

There are at least two other ways to address this argument. One is to say that although phenomenal tokens are ‘intrinsically conscious’, this does not require that every subject they belong to is intrinsic to them. ‘Intrinsically conscious’ might instead simply mean that for the token to exist, there must be at least one subject who undergoes it. Or it might mean that one or more of its subjects is essential to it, but not all: for instance, if a given phenomenal token is associated with some identifiable brain region or system, it might be that this region or system by itself (the token’s minimal supervenience base) is essential to the token, but that other subjects (perhaps the larger systems that contain it, like the brain, head, or whole organism) are not.

The other way is to embrace tight modal connections among phenomenal tokens and their subjects - all of them. Perhaps a given phenomenal token that belongs simultaneously to me, my head, my brain, and a particular neural subsystem is tied to them all essentially: if it did not belong to them all, it would not be the particular token that it is.⁴ Neither of these options seems to us obviously unworkable or problematic, and so we do not feel that there is a successful objection to phenomenal sharing here.

⁴ In section 2 we noted that if there are phenomenal particulars individuated in part by the subject who has them (‘phenomenal havings’), they cannot be shared, by definition. The suggestion made here, that phenomenal tokens are shareable but are individuated in part by the whole *set* of subjects who have them, could be seen as a way to make phenomenal havings, or something very like them, shareable.

Subsection 3.3: The Identity Argument

The next three arguments require us to distinguish between *total* and *partial* phenomenal sharing: between cases where two subjects have exactly the same set of phenomenal tokens (what we might call ‘sharing a phenomenal field’, or ‘sharing a stream of consciousness’), and cases where some phenomenal tokens are shared, but one or both subjects also have some phenomenal tokens that the other does not share.

What we will call the identity argument aims to rule out total sharing specifically, and it does so very directly, by appeal to the idea that subjects are individuated, at a given time, by their total set of experiences - often associated with the stronger claim that subjects are constituted by, identical to, or otherwise nothing over and above a ‘bundle of experiences’. On a view like this, total sharing immediately becomes impossible: if any two subjects shared all of their experiences, they would be one and the same subject. Total sharing would imply identity.

One response would be to deny the conception of subjects that drives this argument - to suppose, for example, that subjects are things like physical organisms, individuated by factors outside their experiences. Two subjects that shared a phenomenal field might then be distinct in virtue of their non-phenomenal properties.

Alternatively, we might accept that subjects are wholly individuated by their experiences, and accept the implication that phenomenal sharing can only ever be partial, not total. This need not be a huge concession, for in cases where we might have been tempted to postulate total sharing, it will likely turn out that on this conception of subjects, the two entities we thought shared a stream of consciousness were never subjects at all, strictly speaking, since they will have the wrong individuation conditions. For example, suppose we had been considering saying that a whole human being and their brain are two subjects who completely share a stream of consciousness. If subjects are wholly individuated by their

experiences, then neither a human being nor a brain can be subjects, since they are (plausibly) individuated at least in part by their physical properties. We would probably think, for instance, that I would cease to exist when my capacity for consciousness is destroyed, but of course my brain may exist after that capability has gone - for a long time, if properly preserved. The experientially-individuated conception of subjects cannot allow that the human and the brain are literally the same subject, since neither is strictly speaking a conscious subject. But what it can say is that these two biological systems are intimately associated in some way with the same subject: they support, or constitute, or generate, or are controlled by, a conscious subject, by supporting, constituting, or generating a set of experiences which themselves are, or constitute, a subject. And then the thought that we wanted to express as ‘the organism and brain are both subjects with the same phenomenal tokens’ would be expressed instead as ‘the organism and brain are both systems which support, constitute, or generate the same subject.’ After all, it is clearly appropriate to say that a given human being has this or that experience - if subjects are wholly individuated by their experiences, this would have to be understood in some not-quite-literal sense of ‘has’. And in that not-quite-literal sense, it might still be perfectly appropriate to say that the brain and the organism ‘share all their experiences’, as long as it is understood that this ‘sharing’ need not mean literally being the subject of those experiences. So we think that the substance of our claims about phenomenal sharing can be sustained whether or not subjects are wholly individuated by their experiences.⁵ Moreover, even if total phenomenal sharing is technically impossible, partial sharing (with the various applications we discuss in section 4) is perfectly possible, because there is nothing especially un-shareable about phenomenal tokens themselves.

⁵ It is also worth noting that on an experientially-individuated conception of subjects, it might well be that Strong Sharing becomes self-contradictory, or collapses into Weak Sharing. For if a subject is, for instance, just a ‘bundle of experiences’, then one natural way to understand ‘mereological overlap’ between two subjects would be that some of the experiences (i.e. phenomenal tokens) in one bundle are also in the other. Phenomenal sharing would be impossible between wholly discrete subjects, because phenomenal sharing is just what it means for two subjects to not be wholly discrete. Since we are defending only weak sharing, this result is no threat to our goals here: indeed it somewhat supports our choice to focus on cases involving mereological overlap.

Subsection 3.4: The **Unity** Argument

While the last argument was against total sharing, the next two arguments focus on partial sharing. The first, which we will call the '**unity** argument', aims to show that partial sharing is impossible, because it would collapse into total sharing (which might either be accepted, or ruled out by the identity argument just considered). The argument can be put as follows:

1. All the phenomenal tokens belonging to a given subject must be phenomenally unified. (The Unity Thesis)
2. A phenomenal token belonging to a certain subject cannot be phenomenally unified with another phenomenal token that does not belong to that subject. (the Boundedness Thesis)
3. *Therefore* If a phenomenal token p were shared between two subjects $S1$ and $S2$, it would be phenomenally unified with all the other phenomenal tokens belonging to $S1$ and also with all the other phenomenal tokens belonging to $S2$. (From 1.)
4. *Therefore* If a phenomenal token p were shared between two subjects $S1$ and $S2$, then $S1$ would have to undergo all the phenomenal tokens belonging to $S2$, and vice versa. (From 3. and 2.)

This argument rests on two premises: the Unity Thesis saying that experiences belonging to the same subject must be unified (defended notably by Bayne and Chalmers 2003, Bayne 2010), and the Boundedness Thesis saying that experiences that are unified must belong to the same subject (critically discussed by Roelofs 2016, 2019, pp.57-72). Both of these are initially plausible, but both might potentially be denied.

Denying the Unity Thesis would mean saying that a single subject might have two or more separate phenomenal fields, not unified with one another. Relying on this to save phenomenal sharing, however, would put severe limitations on how much sharing was possible: it would imply that partial phenomenal sharing required that the shared phenomenal tokens not be unified with the other phenomenal

tokens of at least one of the sharers. Unless I shared my *total* phenomenal field with another subject that overlapped me, the only way for me to share phenomenal tokens with them would be via a rift in either my mind or theirs.

The alternative we prefer is to reject Boundedness, as advocated by Roelofs (2016, 2019 pp.50-53, Roelofs 2019, pp.57-72). This means allowing for ‘between-subjects unity’: even if all the phenomenal tokens of *S1* are phenomenally unified with those of *S2*, it does not follow that *S1* must undergo all the phenomenal tokens of *S2*. It might be that *S1* is conscious of part of *S2*’s experiential field, and that field forms a unified whole, but *S1* is not thereby conscious of the entirety of that whole. Even if there is a unified experiential field that subsumes but goes beyond *S1*’s entire consciousness, *S1* is still only conscious of a part of it. Further exploration of this idea is given in Roelofs 2016, including both canvassing and rebutting arguments that it is self-contradictory, and suggesting a positive phenomenological picture of what it would be like to experience one part of a unified whole. Indeed, one of the objections to allowing between-subjects unity is that, given certain assumptions, it would imply phenomenal sharing - reinforcing the close connections between the two ideas.⁶

A final option might be to relativise unity to a subject, and say that the shared phenomenal token is unified with *S1*’s other phenomenal tokens and disunified with *S2*’s other phenomenal tokens, *relative to S1*, but not relative to *S2* (relative to whom it is unified with *S2*’s other phenomenal tokens, and disunified with *S1*’s). But while such relativity makes sense for functional relations (the role that a given token plays in one mind might be systematically different from the role it plays in another), it is unclear what it would mean for phenomenal unity, mere co-consciousness, to be relative to a subject.⁷ Thus we

⁶ The argument, in brief, is this: we should analyse phenomenal unity as subsumption within a composite experience; thus if one subject has an experience that is unified with an experience of another subject, there will be a single subsuming experience that contains both as parts. On the assumption that experiences require subjects, this suggests a third, overarching, subject must be present, sharing both of the unified experiences with the two original subjects. For more detail see Roelofs 2016, p.3204-3208.

⁷ It is worth noting that the options for between-subjects unity depend significantly on whether or not we accept the claim that phenomenal unity is *transitive* (as argued by Bayne 2010, and denied by Schechter 2014). Transitivity rules out ‘Y-shaped’ patterns of unity, where two internally-unified minds overlap partly but the two unshared parts

feel that the most promising way to defend partial phenomenal sharing from the unity argument is to allow for phenomenal unity between distinct subjects.

Subsection 3.5: The Incompatible Contexts Argument

A final argument against phenomenal sharing turns not just on the unity of consciousness but on its holism, or more precisely on the interdependence of different elements of experience. The sight of someone's lips moving is infused with the sound of their voice; the sweetness of lemonade is sharpened by its sourness; the appearance of one corner pie-shape of the Kanizsa triangle is infused with the influence of the other corners. In all these cases, it might seem artificial to even speak of multiple separate phenomenal tokens, as opposed to a single complex with distinguishable elements. This interdependence is thought to pose a challenge for between-subjects unity, because it implies that the character of the shared phenomenal token will reflect both its unity with the rest of *S1*'s mind, and its unity with the rest of *S2*'s mind, which appears to put it under conflicting demands.

This argument has been made in a few forms and discussed at length in other work, so we will not attempt a thorough treatment (for versions see Basile 2010, 108ff., Coleman 2014, 34ff., and Mørch 2014, 172–175; for responses see Roelofs 2016, pp.3210-3218, 2019, pp.57 ff), but our basic position is that there are two reasonable routes available to addressing it. One is simply to deny the claimed interdependence - to say that while different phenomenal tokens do causally affect one another, perhaps very extensively, they are not interdependent in any non-causal sense (not 'intermingled' or 'infused' into one another, as it is sometimes put) - or at least they are not universally intermingled. Perhaps in some cases they are - perhaps the elements of a gestalt perception cannot exist without one another, and so can only be shared all together, or not at all. But the claim that such an intimate and essential interdependence

are not unified with each other. Cochrane (2021), who defends a phenomenal sharing account of craniopagus conjoinment, denies transitivity in order to preserved the possibility of such an analysis.

is absolutely pervasive in experience is far from obvious, and so denying it is a legitimate option (see Cochrane 2021).

The other option is to accept a really intimate, more-than-causal, interdependence among all unified phenomenal tokens, but say that a subject can have one phenomenal token which is, in this way, intimately shaped by another phenomenal token, even if that subject does not have that other phenomenal token. *S1* need not have all of *S2*'s phenomenal tokens in order to have the one shared token, even though the phenomenal character of that shared token is essentially intertwined with that of *S2*'s other tokens. *S1* might see moving lips *as* infused-with-a-voice, without hearing the voice, or taste sweetness *as* sharpened-by-sourness, without tasting the sourness, or see a pie-shape *as* one corner of a Kanizsa triangle, without seeing the rest of it. Just as we can see, or hear, or read, one part of an integrated unity without grasping all of the other aspects that it essentially depends on, perhaps we can share one element of someone else's unified consciousness without grasping all the other elements that it essentially depends on.

Which of these options is preferable depends on how exactly to understand the interdependence among phenomenal tokens, which is a tricky subject (see, e.g., Dainton 2010, Bayne 2010, Chudnoff 2013). Our point is that for this objection to phenomenal sharing to succeed, this interdependence has to be understood in a very specific way. It has to be something that is both essential to each token, so that a subject who has just one of them is still aware of it, and yet also something that a subject cannot be aware of without having both tokens at once. It must involve, in short, something extrinsic to the experience also being essential to it.⁸ This sort of extreme interdependence comes very close to denying that there are such things as distinct phenomenal tokens: if each one is constituted by all of the others, we might as well say that each subject is really having 'one big experience' at any given moment.

⁸ We are grateful to Donnchadh O'Conaill for discussions on this point.

This ‘one-experience’ view is obviously incompatible with partial phenomenal sharing: if each subject has only one phenomenal token, then two subjects who share a phenomenal token must share their entire phenomenal fields. This view does have defenders (e.g. Tye 2003), but it is a minority view, and eminently deniable; defenders of phenomenal sharing can and should reject it.

Section 4: The Metaphysical Significance of Phenomenal Sharing

Why does phenomenal sharing matter? We think that beyond its intrinsic interest, the possibility of phenomenal sharing has important implications in a number of areas - for how to think about hypothetical or actual cases of connected brains, for how to think about the relationship between ourselves and our parts, and for how to think about the metaphysics of consciousness itself.

Subsection 4.1: Neural Telepathy

First, if phenomenal sharing is indeed possible, then it becomes a natural model for thinking about cases where two brains overlap. The closest to an actually existing case seems to be Krista and Tatiana Hogan, two conjoined twins whose thalami are connected by nervous tissue, and who appear to have some ability to ‘introspectively’ report each other’s perceptions (see e.g. Dominus 2010). Philosophers have discussed the possibility that the twins may ‘share experiences’ (Langland-Hassan 2015, Roelofs 2019, p.112n, Cochrane 2021); the implication of our argument in this paper is that it should be considered a live option that they do. The split-brain case, where a division within a brain gives a measure of autonomy to two neural subsystems, might also be considered in this connection. If there really are two minds here (as argued by, e.g., Schechter 2015, 2018), then it seems clear that they are overlapping minds, since the brainstem, thalamus, and other subcortical structures remain closely connected to each hemisphere. So if there are two minds here, it seems very likely that they share many of their phenomenal tokens.

More broadly, our argument implies that it is perfectly feasible to imagine that technological advances might allow two human beings (or two artificial intelligences) to share phenomenal tokens through connections between their respective brains (or relevant organ/component of consciousness). Hirstein, for example, (2012) imagines this sort of technological ‘mindmelding’ as a sort of medical intervention, for cases where being able to directly examine and compare a patient’s perceptions, emotions, or thought processes could provide vital insight into their symptoms and needs. And several recent papers debate the ethical and political desirability of such technologically-aided integration, including the possibility of human society evolving in the direction of an overarching ‘hive mind’ (Danaher and Petersen 2021, Danaher and Nyholm 2021, Lyreskog et al. 2023). If our argument so far succeeds, then there is no principled incoherence in the idea of this sort of ‘telepathy’.

However, there is an important limitation of what our argument shows. By defending Weak Sharing but not Strong Sharing, we have restricted ourselves to cases where the two subjects sharing a phenomenal token are themselves mereologically overlapping: for all we have argued, two completely discrete subjects cannot share a single phenomenal token (after all, it does seem not possible for two discrete objects to simultaneously share a physical token property). This means that for neural or cybernetic connections to enable genuine ‘telepathy’, they must also bring about a (perhaps temporary) overlap of two minds. What exactly this requires will depend on what determines the boundary of a given conscious subject, a question we have not attempted to answer here. It might be that the boundaries of a conscious subject are the boundaries of their brain: then we would have to ask, for any given sort of neural connection, whether it represents the formation of ‘overlapping brains’, or merely a communication channel between two brains. The split-brain case seems to be a clear example of ‘two overlapping brains’ (or ‘brain-like systems’), but matters are much less clear with the thalamic bridge connecting the Hogan twins’ brains. Arguably we do not know enough at present, and perhaps for the

foreseeable future, to judge whether this or any other case is rightly understood as overlap of brains, or as communication between brains.

Alternatively, it might be that the boundaries of a conscious subject extend beyond the brain - perhaps into other parts of the body (e.g. Brooks 1991), or to the body as a whole biological unity (e.g. Thompson 2010), to further out into the environment (e.g. Vold 2015, Chalmers 2019). Perhaps my mind can be ‘extended’ precisely by things like plugging one end of the right kind of futuristic cable into my head, and the other end into yours: perhaps the information transfer this enabled would expand the boundaries of my mind, and simultaneously expand the boundaries of yours, so that some of the neural activity in my brain is realising a phenomenal token that is both part of my mind and part of yours.

Our case for phenomenal sharing is neutral about the boundaries of minds, and so it is only one part of a full theoretical treatment of the kind of ‘neural telepathy’ that might be technologically possible, and which might be actual in split-brain or craniopagus cases. But it is an important part: clearing away objections to phenomenal sharing allows us to do justice to the forms of mental overlap that are already possible and that are likely to become only more prevalent as neural technology develops.

Subsection 4.2: Too Many Minds

A much more familiar form of ‘mental overlap’ is between such entities as a human being, their head, and their brain. Clearly these entities overlap, and intuitively they are conscious, but we do not normally think of them as ‘overlapping conscious subjects’ - indeed, lines of thinking that seem to entail thinking of them that way are often felt to pose a ‘Too Many Minds’ problem. If my brain is conscious, but I am a whole human being (arms, legs, and all) then are there two conscious subjects here, where I thought there was one? Yet if my brain is not conscious, why not - surely it is capable of consciousness: if it were held in a vat on life-support, it could be conscious, so how could the presence of my skull and body somehow stop

it being conscious? Of course these questions may be bypassed by denying that the brain itself is really conscious - perhaps it is only conscious in a derivative sense, through supporting a certain bundle of experiences (which are also supported by the whole human being). But those who think of at least some physical systems as themselves conscious subjects need to address the likelihood that suitable physical systems may often overlap with each other (Blackmon, 2016, 2021, argues this at length for the two cerebral hemispheres).

We have already mentioned Sutton's (2014) approach to this problem: allowing multiple distinct but overlapping beings to have 'the same mind', so that although there are many thinkers, "there is not too much thinking" (p.619) because they share a supervenience base for their thoughts. Our arguments here aim to reinforce this approach and to show its application to other issues. More broadly, phenomenal sharing is a necessary part of any 'manyist' solution, i.e. one which accepts that there are many conscious beings but argues that their multiplicity is innocuous or unobjectionable (e.g. Lewis 1993). To many, it has seemed that a multiplicity of overlapping physical objects is acceptable, but a multiplicity of overlapping conscious subjects is not (e.g. Unger 2004, Merricks 2005): one possible explanation of this difference is that property-sharing is usually assumed possible for physical objects but not for minds. (For a more detailed discussion of how manyism relates to phenomenal sharing, see Roelofs 2022).

Subsection 4.3: The Varieties of Panpsychism

Finally, the possibility of phenomenal sharing impacts the prospects for certain theories of consciousness, in particular panpsychism. Indeed, the most explicit recent discussions of experience-sharing have occurred in relation to panpsychism (see, e.g., Basile 2010, p.108, Roelofs 2016, pp.3205-3209, 2019, 2020-a, pp.61-69, Miller 2017, p.12). According to panpsychists, humans are conscious because consciousness is inherent in matter; our experiences arise somehow from the simpler experiences of the

physical particles that make us up. (Depending on the version of panpsychism, intermediate-sized entities, like cells or neural assemblies, may also have consciousness, built up from that of smaller things and feeding into that of larger things.) But a lot depends on how this ‘arising from’ is understood. Some panpsychists take human minds to be ‘strongly emergent’ phenomena, so that macro-level subjects (despite having parts) are fundamental entities in their own right, and their phenomenal properties are fundamental properties. Other panpsychists, whose view might be called ‘constitutive’ or ‘weakly emergentist’, take macro-level subjects to be fully explained by the simpler subjects that compose them, just as macro-level physical objects are fully explained by the simpler objects that compose them.

Accepting or denying phenomenal sharing does not decisively establish or refute any version of panpsychism, but it is extremely helpful to panpsychists in proportion as they adopt a constitutive view. To illustrate, we will briefly sketch two versions of panpsychism that centre on phenomenal sharing: one more constitutive/weakly emergentist, preferred by Roelofs, and one involving a hybrid of constitutive/weakly emergentist and strongly emergentist elements, preferred by Goff.

Subsection 4.4: Constitutive Panpsychism

Phenomenal sharing offers multiple advantages to constitutive panpsychists, all stemming from the close parallel it enables between mental combination and physical combination. Physical parts together form innumerable physical wholes at different levels, without explanatory gaps, causal competition, or unparsimoniousness. And a key part of why they can do that is that the wholes do not add anything wholly new - no wholly new powers, qualities, or other features - but only organise and arrange what is already there. In short, it seems like physical combination is constitutive (i.e. involves only weak emergence) *because* physical wholes share the properties of their parts, with any properties not shared with their parts coming from the way that the properties that are shared get put together.

The ambition of **constitutive** panpsychism is that mental combination might follow the same model: the properties of complex mental wholes (like our minds) might be simply a new arrangement and organisation of the properties already present at the fundamental level, rather than brute ontological novelties. The fundamental parts of me have very simple experiences; larger and more complex parts share these experiences, but combine them in novel ways based on their structure, and so on up to me myself, who shares all these experiences (or some subset of them, depending where we draw the boundaries of ‘me’), but for whom they form a complex structure not present in any of them individually.

If that makes sense, mental combination could potentially resemble physical combination also in its lack of explanatory gaps, causal non-competition, and multiplicity without unparsimoniousness. Of course, actually making good on this potential is far from easy: **constitutive** panpsychists need to address the various ‘combination problems’ that have been raised by critics (see Seager 1995, 2010 Goff 2009-a, 2009-b, 2015, 2019, forthcoming, Coleman 2014, Mørch 2014, Roelofs 2014, 2016, 2019, 2020-a, 2020-b, Chalmers 2017). But without phenomenal sharing, the effort seems almost dead in the water: analysing the new properties of whole as complex arrangements of phenomenal properties shared with their parts is doomed to failure if they cannot even share the simplest phenomenal properties of those parts. It would be like trying to explain the macrophysical properties of a desk on the assumption that for it to have mass, charge, or any other physical causal powers, it must have them *in addition to* the powers of its parts. Phenomenal sharing does not guarantee that **constitutive** panpsychists will successfully solve the combination problems, but it is a vitally useful premise for their attempts (for discussion see in particular Roelofs 2019, 2020-a).

Subsection 4.5: Hybrid Panpsychism

Phenomenal sharing is also compatible with certain sorts of qualified emergentism: more precisely, it is compatible with thinking that while every *subject* is fundamental, the experiences of some subjects are shared with, or grounded in, the phenomenal properties of other subjects. According to *hybrid panpsychism*:

- there are fundamental laws ensuring that when conscious particles are arranged in certain combinations, a subject associated with the system as a whole emerges,
- these emergent subjects do not come into existence with new experiences; rather, they share experiences with, or have experiences that are grounded in the experiences of, subjects at more basic levels of nature.

We can further develop this idea by exploring how it helps avoid some of the ‘combination problems’ that have been raised for weakly emergentist/constitutive panpsychism, in particular the subject-summing argument (cf. Goff 2009a, 2009b, Roelofs 2020-a) and the revelation argument (cf. Goff 2015, 2017-b: ch. 5, Roelofs 2020-b).

We can pose the subject-summing problem as a kind of conceivability argument:

1. For any group of micro-level subjects standing in certain physical relationships to each other, we can conceive of a world in which those micro-level subjects stand in those physical relationships without it being the case that they compose a macro-level subject.
2. If 1, then, for any state of affairs involving a number of micro-level subjects standing in certain physical relationships to each other, that state of affairs does not necessitate the existence of a subject non-identical with any of the micro-level subjects.

3. If 1 and 2, then, for any state of affairs involving a number of micro-level subjects standing in certain physical relationships to each other, that state of affairs does not ground the existence of a macro-level subject.
4. Conclusion: We cannot ground macro-level subjects in facts about micro-level subjects bearing physical relations to each other.

The inference from (1) to (2) involves a move from conceivability to possibility. It is commonly supposed that there are gaps between conceivability and possibility, e.g. we cannot infer from the conceivability to the possibility of water existing without H₂O. However, this kind of gap between conceivability and possibility is brought about by the fact that the concept ‘water’ does not pick out its referent in terms of its essential nature. The properties and kinds referred to in (1), in contrast, *are* conceived of in terms of their essential nature – see discussion of the revelation argument shortly to follow – which makes plausible the inference from conceivability to possibility asserted in (2). The inference from (2) to (3) involves a broadly agreed principle that if fact X grounds fact Y, then necessarily if X then Y.

Weakly emergentist panpsychists must meet this argument head on, either by denying premise 1 (see Roelofs 2019, p.105 ff) or by denying premise 2 (see Goff 2017-a). But the qualified emergentism under consideration here has a straightforward way around this problem: subjects are *not* grounded in any more basic phenomenal facts, even though their experiences *are*. How exactly the experiences of emergent subjects are derived from the experiences of more basic subjects will become clearer as we show how hybrid panpsychism can respond to the revelation argument.

The revelation argument starts from the ‘revelation thesis’ (Goff 2015, 2017: ch. 5), the idea that not only that we are aware *of* our mental states, in having them, but that in such awareness we grasp their essential nature. Contrast this with my awareness of the hum of my computer. The thing I am aware of – the humming – is constituted of vibrations in the air. But in my awareness of it, I do not conceive of it as

such. The pain I am currently experiencing, on the other hand, is a feeling, and in my awareness of my pain I conceive of it as such. To have that pain just is to feel a certain way; and I am aware, in my awareness of my pain, of that precise way to feel. It is arguably self-evident, upon reflection, that in attending to our feelings and thinking of them in terms of how they feel, we grasp their essential nature. But this seems to yield the following argument against weakly emergentist panpsychism (cf. Chalmers 2017, p.190):

1. *Revelation* – In my awareness of my pain, its nature is revealed to me.
2. If trillions of micro-level phenomenal properties constitute my pain, then this is part of its nature.
3. It's not the case that, in my awareness of my pain, its constitution by trillions of micro-level phenomenal properties is revealed to me.
4. It's not the case that my pain is constituted of trillions of micro-level phenomenal properties.

Weakly emergentist panpsychists who accept the revelation thesis must either deny premise 2 (and claim that something can constitute my pain but not be part of its nature) or deny premise 3 (and claim that introspection does in fact reveal a vast array of microexperiences). Roelofs (2020-b) argues that there are different senses of 'revelation' at play here, yielding importantly different versions of the revelation argument - in some, the counterpart to premise 2 is false, while in others the counterpart to premise 3 is false. But defending this line is admittedly subtle: one must thread a conceptual needle. It might be an advantage, then, that hybrid panpsychism can potentially side-step the whole argument, just as it does with the subject-summing argument.

Hybrid panpsychism does this by appealing to the two-step hypothesis. Consider a slight pain I am feeling at a given moment – call it 'P' – and the region of my brain whose activity is correlated with that pain – call it 'R'. If panpsychism is true, then R is ultimately constituted of trillions of conscious micro-level particles. Given this, it is extremely plausible that, if R instantiates macro-level consciousness

at all, it will instantiate an extremely complex conscious state: call that state ‘the busy state’. But, so long as we are comfortable with mental overlap, we needn’t hold that the busy state is identical with P, or indeed with any of my conscious states. We might instead take P to be a *partial abstraction* from the busy state, drawing inspiration from David Armstrong’s (1997) view of states of affairs. For Armstrong, states of affairs - e.g. the state of affairs of the-electron-being-negatively-charged - are fundamental. We reach a conception of universals by taking a given state of affairs – or ‘thick particular’ – and *partially considering it*. In other words, we abstract away from certain aspects of the whole state of affairs. To illustrate this with the above example, I partially consider the state of affairs of the-electron-being-negatively-charged in order to reach a conception of the universal of negative charge. Despite taking them to be non-fundamental, it is clear that Armstrong does not take universals (or substrata for that matter) to be merely shadows of language: our talk of universals carves nature at the joints. Indeed, given that both the state of affairs and its universals are natural entities, it seems to us that we could think of the relation between them as a form of grounding. The universals are grounded in the states of affairs in virtue of being subsumed within their being; we call this form of grounding ‘grounding by subsumption.’

On the **hybrid** view the relationship between the busy state and P (my pain) is also a case of grounding by subsumption. P is a perfectly natural property, but is also a partial abstraction from another perfectly natural property, namely the busy state. The busy state grounds P by subsuming it within itself.

Trillions of micro-level phenomenal properties do contribute indirectly to P, in virtue of constituting the state which P is a partial abstraction from (namely, the busy state). However, given that the nature of P abstracts away from a lot of the detail of the busy state, it is to be expected that I am not aware of the trillions of micro-level phenomenal properties that constitute the busy state. In other words, there is a two-stage chain of constitution from certain micro-level phenomenal properties to my

experiences, and the second stage of the constitution (grounding by subsumption) explains why I am not aware of so many of the micro-level properties which, strictly speaking, constitute my experience.

Bringing in the emergence of subjects, we can summarise this more developed form of hybrid panpsychism as follows:

- *Step 1* – It is a basic law of nature that when micro-level subjects, $M_1, M_2 \dots M_n$, stand in certain physical relations to another, the resulting state of affairs causes a fundamental subject S to emerge, such that: (i) S is composed of all and only $M_1, M_2 \dots M_n$, and (ii) S shares all and only the phenomenal properties of $M_1, M_2 \dots M_n$. Call such a subject a ‘basic macro-level subject’.
- *Step 2* – It is a basic law of nature that when a basic macro-level subject emerges, it causes numerous other co-located subjects to emerge, such that the phenomenal properties of those subjects are grounded by subsumption in the phenomenal properties of the basic macro-level subject. (Obviously both principles leave out a lot of detail that would need to be filled in on the basis of empirical investigation).

On this view, the mind of a human or an animal is a non-basic macro-level subject, whose experiences are grounded by subsumption in the experience of the basic subject with which it is co-located.

Despite these advantages, emergentism about subjects does come with a cost: each of the multiplicity of radically emergent subjects is an ontological addition, and thus ‘costs extra’ when reckoning parsimony: hybrid panpsychism is thus less parsimonious than weakly emergentist panpsychism. However, the principles which produce these many subjects are simple and elegant, and they arrange themselves neatly around a shared set of phenomenal property instances. The cost, although real, is negligible.

Strong emergentism about subjects might also be thought to raise causal exclusion worries: assuming the micro-level is causally closed, emergent macro-level subjects will be either epiphenomenal or have their causal effects over-determined. However, there is no causal exclusion problem here so long as we take the basic relation of causation to be phenomenal tokens rather than phenomenal havings or subjects. What we want to account for is the fact that my *phenomenal properties*, e.g. my pain, my desire, my cognitive states, etc. have causal influence in the world. It's not clear that there is a further obligation to account for the causal efficacy of my *mind* in addition to the causal efficacy of its states.

For these reasons, hybrid panpsychism is arguably able to avoid both the problems of weak emergence and those of strong emergence. This makes it a view worth taking very seriously indeed.

Conclusions

It's common to think of experiences – or more precisely, token instances of phenomenal properties – as essentially exclusive, each belonging to only one subject and never shared with any others. Thinking of them this way makes them sharply different from physical properties. But in order to properly understand the mind's place in nature we have to be open to re-examining our assumptions about the mind, especially assumptions about radical differences between mind and matter. We have tried to show that if the possibility of phenomenal sharing is approached with an open mind, we will not find insuperable conceptual objections, and may find new insights or hermeneutical tools for understanding our minds, other minds, and how they might be related.

References

Armstrong, David. (1997) *A World of States of Affairs*, Cambridge University Press.

Basile, Pierfrancesco. (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.

Bayne, T. (2010). *The Unity of Consciousness*. Oxford: Oxford University Press.

Bayne, T., and Chalmers, D. (2003). “What Is the Unity of Consciousness?” in A. Cleeremans (ed.), *The Unity of Consciousness: Binding, Integration, Dissociation*. Oxford: Oxford University Press.

Blackmon, James. (2016). “Hemispherectomies and Independently Conscious Brain Regions”, *Journal of Cognition and Neuroethics* 3/4: 1-26.

Blackmon, James (2021). “Integrated Information Theory, Intrinsicity, and Overlapping Conscious Systems”, *Journal of Consciousness Studies* 28/11–12: 31–53.

Brooks, R., 1991, “Intelligence without representation,” *Artificial Intelligence*, 47: 139–159.

Brüntrup, Godehard, and Jaskolla, Ludwig (Eds.). (2017). *Panpsychism: Contemporary Perspectives*. Oxford University Press.

Buchanan, Jed, and Roelofs, Luke. (2019). “Panpsychism, Intuitions, and the Great Chain of Being.” *Philosophical Studies* 176 (11): 2991-3017.

Burke, Michael (1994). “Dion and Theon: An Essentialist Solution to an Ancient Puzzle.” *The Journal of Philosophy* 91 (3): 129-139.

Chalmers, David. (2017). “The Combination Problem for Panpsychism.” In G. Brüntrup and L. Jaskolla (Eds.) *Panpsychism: Contemporary Perspectives*. Oxford University Press: 179-214.

Chalmers, David. (2019). “Extended Cognition and Extended Consciousness.” In Matteo Colombo, Elizabeth Irvine & Mog Stapleton (eds.), *Andy Clark and his Critics*. Wiley-Blackwell: 9-20.

Chudnoff, E. (2013). “Gurwitsch’s Phenomenal Holism.” *Phenomenology and the Cognitive Sciences* 12 (3): 559–578.

Cochrane, Tom. (2021). A Case of Shared Consciousness. *Synthese* 199, pages 1019–1037

Coleman, S. (2014). “The Real Combination Problem: Panpsychism, Micro-Subjects, and Emergence.” *Erkenntnis* 79 (1): 19–44.

Dainton, B. (2010). "Phenomenal Holism." *Royal Institute of Philosophy Supplement* 67: 113–139.

Danaher, J., and Nyholm, S. (2021). "Should we use technology to merge minds?" *Cambridge Quarterly of Healthcare Ethics* 30 (4): 585–603.

Danaher, J., and Petersen, S. (2021). "In Defence of the Hivemind Society." *Neuroethics* 14 (2): 253–267.

Dominus, S. (2011). "Could Conjoined Twins Share a Mind?" *New York Times Magazine*, May 29. <http://www.nytimes.com/2011/05/29/magazine/could-conjoined-twins-share-a-mind.html>.

Ehring, D. (1996). "Causation, Determinables, and Property Instances." *Nous* 30 (4): 461–480.

Geach, Peter. (1967). "Identity." *Review of Metaphysics* 21: 3-12.

Goff, P. (2009a). "Can the Panpsychist Get Round the Combination Problem? In D. Skrbina (ed.), *Mind That Abides: Panpsychism in the New Millennium*. Amsterdam: John Benjamins.

Goff, P. (2009b). "Why Panpsychism Doesn't Help Us Explain Consciousness." *Dialectica* 63 (3): 289–311.

Goff, Philip. (2013). "Orthodox property dualism + linguistic theory of vagueness = panpsychism." In R. Brown (Ed.) *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*, Springer: 75-91.

Goff, Philip. (2015) "Real acquaintance and physicalism," in P. Coates & S. Coleman (Eds.) *Phenomenal Qualities: Sense, Perception, and Consciousness*, New York: Oxford University Press.

Goff, Philip. (2017-a). "The Phenomenal Bonding Solution to the Combination Problem." In G. Brüntrup and L. Jaskolla (Eds.) *Panpsychism: Contemporary Perspectives*. Oxford University Press: 283-303.

Goff, Philip. (2017-b). *Consciousness and Fundamental Reality*. Oxford University Press.

Goff, Philip (2015). 'Against constitutive Russellian monism,' in T. Alter & Y. Nagasawa (Eds.) *Russellian Monism*. Oxford University Press.

Goff, Philip (2019) 'Grounding, analysis and Russellian monism,' in S. Coleman (Ed.) *The Knowledge Argument Then and Now*, Cambridge University Press.

Goff, Philip (Forthcoming) ‘*Is consciousness fundamental?*’ in G. Rabin (Ed.) *Grounding and Consciousness*, Oxford University Press.

Hirstein, William. (2008). “Mindmelding: Connected Brains and the Problem of Consciousness.” *Mensa Sana Monographs*, 6 (1): 110-130.

Hirstein, William. (2012). *Mindmelding : Consciousness, Neuroscience, and the Mind's Privacy*. Oxford University Press.

Langland-Hassan, P. (2015). “Introspective Misidentification.” *Philosophical Studies* 172 (7): 1737–1758.

Lewis, David. (1993). “Many, but Almost One.” In J. Bacon, K. Campbell, and L. Reinhardt (Eds.) *Ontology, Causality, and Mind: Essays in Honour of D.M. Armstrong*, Cambridge University Press: 23-45.

Lyreskog, David, Zohny, Hazem, Savulescu, Julian, and Singh, Ilina. (2023). “Merging Minds: The Conceptual and Ethical Impacts of Emerging Technologies for Collective Minds.” *Neuroethics* 16:12 <https://doi.org/10.1007/s12152-023-09516-3>

Merricks, Trenton. (2005). “Composition and Vagueness.” *Mind* 114: 615-637.

Mørch, Hedda Hassel. (2014). *Panpsychism and Causation: A New Argument and a Solution to the Combination Problem*. (Doctoral Dissertation) Department of Philosophy, Classics, History of Art and Ideas, University of Oslo, Oslo.

Miller, Greg. (2017). “Can Subjects Be Proper Parts of Subjects? The De-combination Problem.” *Ratio* 30 (2): 1–18.

Mørch, Hedda Hassel. (2019a). “Is Consciousness Intrinsic?: A Problem for the Integrated Information Theory.” *Journal of Consciousness Studies* 26 (1-2):133-162.

Mørch, Hedda Hassel. (2019b). “Is the Integrated Information Theory of Consciousness Compatible with Russellian Panpsychism?” *Erkenntnis* 84 (5): 1065-1085.

Noonan, Harold. (2003). *Personal Identity*, 2nd ed. First edition 1989. Routledge

Noonan, Harold. (2010). “The Thinking Animal Problem and Personal Pronoun Revisionism.” *Analysis* 70:1, pp.93-98

- Oizumi, Masafumi, Albantakis, Larissa, and Tononi, Giulio. (2014). "From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0." *PLOS Computational Biology* 10 (5). Online.
- Olson, Eric. (2003). "An Argument for Animalism." In R. Martin and J. Barresi (Eds.), *Personal Identity*, Blackwell: 318-335.
- Putnam, Hilary. (2003). "The Nature of Mental States" (1965). In T. O'Connor, D. Robb, and J. Heil (eds.), *Philosophy of Mind: Contemporary Readings*. New York: Routledge.
- Roelofs, Luke. (2016). "The Unity of Consciousness, within and between Subjects." *Philosophical Studies* 173 (12): 3199–3221.
- Roelofs, Luke. (2019). *Combining Minds: How to Think about Composite Subjectivity*. Oxford University Press.
- Roelofs, Luke. (2020-a). "Can We Sum Subjects? Evaluating Panpsychism's Hard Problem." In W. Seager (ed.), *The Routledge Handbook of Panpsychism*. New York: Routledge: 245-258.
- Roelofs, Luke. (2020-b). "Consciousness, Revelation, and Confusion: are constitutive panpsychists hoist by their own petard?" *Dialectica* 74 (1): 59-90.
- Roelofs, Luke. (2022). "No Such Thing As Too Many Minds." *Australasian Journal of Philosophy*. <https://doi.org/10.1080/00048402.2022.2084758>
- Rosenthal, D. (1986). "Two Concepts of Consciousness." *Philosophical Studies* 49: 329–359.
- Rosenthal, D. (2005). *Consciousness and Mind*. Oxford: Oxford University Press.
- Schechter, E. (2014). Partial Unity of Consciousness: A Preliminary Defense. In D. Bennett, D. J. Bennett & C. Hill (eds.) *Sensory integration and the unity of consciousness*, Cambridge, MA: MIT Press, 347-374.
- Schechter, E. (2015). "The Subject in Neuropsychology: Individuating Minds in the Split-Brain Case." *Mind and Language* 30 (5): 501–525.
- Schechter, E. (2018). *Self-Consciousness and "Split" Brains: The Minds' I*. Oxford: Oxford University Press.

- Seager, William. (1995). "Consciousness, Information and Panpsychism." *Journal of Consciousness Studies* 2–3: 272–288.
- Seager, William. (2010). "Panpsychism, Aggregation and Combinatorial Infusion." *Mind & Matter* 8 (2): 167-184.
- Seager, W. (2017). "Panpsychist Infusion." In G. Brüntrup and L. Jaskolla (eds.), *Panpsychism: Contemporary Perspectives*. Oxford: Oxford University Press.
- Seager, William (Ed.). (2019). *The Routledge Handbook of Panpsychism*. Routledge.
- Shoemaker, Sidney. (1997). "Parfit on Identity." In J. Dancy (Ed.), *Reading Parfit*, Malden, MA: Blackwell.
- Sider, Ted. (2003). "Maximality and Microphysical Supervenience." *Philosophy and Phenomenological Research* 66: 139–149.
- Simon, Jonathan. (2017). "The Hard Problem of the Many." *Philosophical Perspectives* 31 (1): 449-468.
- Skrbina, David. 2005. *Panpsychism in the West*. MIT Press.
- Sutton, C. S. (2014). "The Supervenience Solution to the Too-Many-Thinkers Problem." *Philosophical Quarterly* 64 (257): 619–639.
- Thompson, Evan, (2010). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.
- Unger, Peter. (2004). "Mental Problems of the Many." In. D. Zimmerman (Ed.) *Oxford Studies in Metaphysics*, Volume 1, Clarendon Press: 195-222.
- Vold, Karina. (2015). "The parity argument for extended consciousness." *Journal of Consciousness Studies* 22: 16-33.