

No Such Thing as Too Many Minds

Luke Roelofs

To cite this article: Luke Roelofs (2022): No Such Thing as Too Many Minds, Australasian Journal of Philosophy, DOI: [10.1080/00048402.2022.2084758](https://doi.org/10.1080/00048402.2022.2084758)

To link to this article: <https://doi.org/10.1080/00048402.2022.2084758>



Published online: 10 Jul 2022.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)



No Such Thing as Too Many Minds

Luke Roelofs

New York University

ABSTRACT

Many philosophical views have the surprising implication that, within the boundaries of each human being, there is not just one mind, but many: anywhere from two (the person and their brain, or the person and their body) to trillions (each of the nearly-entirely-overlapping precise entities generated by the Problem of the Many). This is often treated as absurd, a problem of 'Too Many Minds', which we must find ways to avoid. It is often thought *specifically* absurd to allow such a multiplication of conscious subjects, even if we could accept it for physical objects. I consider metaphysical, phenomenological, and moral arguments for this asymmetry, and show that they all fail: many overlapping conscious minds is no more problematic than many overlapping physical objects. Theories that imply such a multiplicity may or may not be true, but they cannot be rejected simply for implying it.

ARTICLE HISTORY Received 30 March 2021; Revised 8 May 2022

KEYWORDS Consciousness; Manyism; Problem of the Many; Mereology; Overlap; Parthood

1 Introduction

It can be dispiriting to write papers only a few people read. So it is comforting that, according to some philosophers, even the most neglected paper will be read by several trillion numerically distinct people, for there are trillions of people wherever we thought there was only one. Whatever succour this idea might give vain writers, most who have considered it have thought it absurd, and accordingly rejected any theory which entails such a result. I disagree: there is nothing implausible about there being, for 'each of us', multiple overlapping conscious subjects in one spot.

Section 2 enumerates various forms of 'Too-Many-Minds' problem, all suggesting that where we thought we saw a single minded being, there are really many. Versions of this problem arise concerning people and their bodies, bodies and their heads, people and all-of-them-but-the-left-foot, and even people-precisified-as-including-this-molecule and people-precisified-as-excluding-this-molecule. And recent debates about panpsychism raise additional Too-Many-Minds worries.

Many different views imply such multiplicity: I won't evaluate any of them, just argue that this implication gives no reason to reject them. Similarly, where Too-Many-Minds seems to follow from common-sense naturalist assumptions, I won't judge whether these assumptions are true, or really have this implication: I just argue that generating many minds gives no reason to reject them. This acceptance of many minds has come to be called 'manyism'.

I distinguish metaphysical, phenomenological, and moral objections to manyism about conscious subjects. Metaphysical objections focus on *parsimony* and *epiphenomenalism*; section 3 responds that wholes neither add to their parts when gauging parsimony, nor compete with their parts for causal efficacy. This solution succeeds, I think, when applied to non-conscious things. But extending it to conscious subjects requires token conscious experiences to be shareable between distinct (but not disjoint) subjects, which challenges traditional ideas about mental privacy. In other work I defend experience-sharing more fully [Roelofs 2019; Goff and Roelofs forthcoming]; section 3 offers a preliminary defence and shows how, *if* overlapping subjects can share experiences, the metaphysical challenges to manyism fail.

However, experience-sharing adds bite to phenomenological objections, which hold that if manyism were true, it would be manifest in our experience in certain ways that it is not. Section 4 considers and rebuts two phenomenological objections: that an experience's owners should be revealed by reflection on it, and that our consciousness should show internal boundaries which it does not show.

A final objection is *moral*: if manyism were true, simple moral principles like 'prefer helping more beings to helping fewer' would seem impossible to apply or would have absurd consequences. Section 5 responds that our moral intuitions have developed through our interactions with each other, and so *if* manyism is true, they have developed through our interactions with clouds of trillions of overlapping subjects. Given this, they should be interpreted in manyism-friendly ways.

I conclude that there are no good metaphysical, phenomenological, or moral reasons to reject manyism about minds. The idea that each of us is really a huge cloud of overlapping conscious subjects may be like the idea that our mental lives depend on a wrinkly piece of meat, or that some of our ancestors were monkeys: a surprising, and in some ways counter-intuitive, implication of naturalism that we should work to accept rather than seek ways to avoid.

2 How Many Minds is Too Many?

Philosophers have offered four sorts of argument, appealing to different sorts of part, that any situation we would count as containing 'one person' actually contains multiple overlapping persons.¹ Let us briefly survey these four types.

2.1 Four Routes to Too Many Minds

The first sort of Too-Many-Minds problem involves familiar overlapping things, like a human being and their brain. Most people think that brains exist, but Merricks claims that admitting the existence of brains leads to an 'unacceptable multiplication of thinkers':

Within the region occupied by a human organism there is a conscious human organism and a conscious brain. The brain is not identical with the organism; they differ in properties and parts. So ... there are at least two conscious entities within that region ... [Merricks 2001: 49]

¹ For discussion see Unger [1980, 2004], Lewis [1993], Merricks [2001], Sider [2003], Seager [2010], Zimmerman [2010], Blackmon [2016, 2021], Simon [2017], Mørch [2019].

To avoid this result, Merricks denies the existence of brains; to follow through he also denies the existence of heads, skulls, brainstems, etc.

The second sort of multiplication involves any part of me that includes my brain, such as arbitrary sections ('my top half') and complements of other parts ('all of me except my left foot'). Plausibly, I could be conscious even if I were reduced to nothing more than such a part (by losing my left foot). And so, the argument goes, if 'all of me except my left foot' currently exists, it is a conscious subject, distinct from me.

The third Too-Many-Minds problem, sometimes called 'the Problem of the Many' [Unger 1980], involves entities differing from me only in my 'questionable' microscopic parts. For instance, consider a water molecule in the process of leaving one of my skin cells, neither determinately still in nor determinately out yet. Now consider the entity which contains all of me except this molecule, and the entity which contains this molecule as well. Both seem to be conscious thinkers, and 'equally strong candidates' for being *me*—there is no non-arbitrary reason to privilege one as 'the whole person'.

A final sort of Too-Many-Minds problem arises for panpsychist theories, which posit consciousness more widely than we normally think. This potentially means consciousness in parts of me which do not include my whole brain: if neurons, or electrons, or even the whole universe, are themselves conscious, it looks like there are *way* too many minds.

All four arguments focus on 'parts' in a certain core sense: parts of something that exist at the same time as it and belong to the same ontological category. Parallel arguments could be raised in terms of temporal parts, like the 'part' of me that existed from 2016 to 2017 [Lewis, 1976; Noonan 2003; Olson 2003], or logical parts, like my property of being rational considered apart from my property of being human [Paul 2002, 2006]. Parallel arguments might also be made about particular mental states, such as an argument that functionalism implies that each mental state has a multiplicity of overlapping neural realisers. In particular, Johnston [2017a, 2017b, 2021] makes an argument paralleling Simon's moral objection to manyism: naturalism seemingly implies the existence of myriad 'personites', person-like beings existing for only part of a person's lifespan, and this renders ethical reasoning intractable.

My focus is on addressing objections to multiplicity involving synchronic physical parts of subjects being themselves subjects. Considering the multiplication of conscious temporal parts, logical parts, etc. would introduce issues, such as those concerning persistence over time, that are beyond the scope of this paper. But to the extent that analogous issues arise in both cases, I think my arguments may be more broadly applicable.

All four sorts of Too-Many-Minds problem have a similar structure, with three basic premises.

- (1) The entities in question are intrinsically suitable to be conscious.
- (2) Consciousness is intrinsic: whether something is conscious does not depend constitutively on anything outside that thing.
- (3) The effects of the other parts (or other continuers) on the entities in question are not such as to harm, suppress, or disrupt their consciousness.

The first premise is often supported by observing that the parts in question would plausibly remain conscious if the other parts were removed, as long as the right external support were provided (for instance, my head could remain conscious without the rest of me, if supplied with nutrients, stimulation, etc.). The second and third premises are intuitive; it seems that whether I'm conscious shouldn't depend constitutively on something outside of me, and it would be bizarre to think of my neck as somehow undermining or disrupting the consciousness of my head. But if all three premises are true, it seems to follow that the parts in question are presently conscious.

Of course the plausibility of these premises may differ among the different versions of the argument: someone might accept that my brain is capable of consciousness, but deny that any subsystem of it is, let alone assemblages of putatively conscious atoms. We might even disagree about which of these entities exist: top halves of people might not deserve a place in our ontology. In this paper I am neutral about which things are conscious and which things exist: I aim to refute a particular kind of argument, that certain things *must not* exist or *must not* be conscious, on pain of generating too many minds. No number of minds, I argue, is too many.

2.2 Accepting and Resisting Responses

I prefer to accept the profusion of minds and reconcile myself to it. An 'accepting' response to this multiplicity of minds aims to make multiplicity palatable—saying that *yes*, there are that many minds, but *no*, that isn't too many. After all, a multiplicity of material objects is not so bad, and is often admitted by the same writers who reject multiplying minds [Merricks 2001: 49, 95, 106; Unger 2004: 378–79]. A common label for accepting responses is 'manyism': 'solutions ... according to which every candidate is an experiencer' [Simon 2017: 451].²

Many philosophers have found this profusion of minds absurd, and have looked for ways to avoid it and to rescue our initial intuitive thought that there is sometimes just one conscious subject in a given place. Call this style of response 'resisting'.

There are various forms of resisting response. Unger, for instance, sees the Problem of the Many as threatening enough that to avoid multiplication of minds he has at different times denied the existence of people, including himself [1979], or endorsed substance dualism [2004]. Zimmerman [2010] and Simon [2017] also use the threat of multiplication as an argument for substance dualism [cf. Johnston 2017a: 641, 2017b: 227]. Merricks exhibits a more subtle sort of resisting response: he remains technically a substance monist, but in order to avoid admitting that 'my thoughts are not mine alone, but shared' [2001: 49], Merricks denies the existence of brains, other internal organs, arbitrary sections, and all but one of the all-but-indiscernible entities involved in the Problem of the Many. All of these solutions posit non-physical facts singling out a unique 'me' that cannot be inferred from even the fullest physical knowledge.

A different style of resisting response comes from Tononi, who imposes an 'exclusion postulate' by which consciousness on one level precludes consciousness at other levels [Tononi 2012: 59–68; Tononi and Koch 2015: 6]. Others adopt a 'fusionist'

² Blackmon [2016, 2021] defends the likely reality of multiple overlapping minds for each human being, though he does not use the term 'manyism', and leaves open the scale of the multiplicity—he mostly focuses on the overlapping minds of hemispheres and whole brains.

account of mental combination, on which conscious parts vanish when they form a conscious whole [Seager 2010, 2017; Mørch 2019]. Both allow the consciousness of a physical system to come apart from its physical properties: that is why we can say that my left-hemisphere-complement would be conscious if the left hemisphere were removed, but not as long as it remains attached. Each part deprives the others of consciousness.

In this paper I will not try to motivate manyism, though I do think it is a much more elegant and attractive position than the contortions required for a resisting response. Rather, I consider objections which claim that, while we might accept a profusion of overlapping physical objects, there is something special about *consciousness* which rules out a profusion of conscious minds. This fits into a broader pattern of philosophers rejecting any sort of mereological overlap of conscious minds, affirming what is sometimes called an ‘Anti-Nesting’ principle [Putnam 1965: 215–16; Kammerer 2015; Schwitzgebel 2015; Mørch 2019: 135]. In defending manyism I will thus be defending the more general possibility of conscious beings overlapping.

3 Metaphysical Arguments Against Manyism

I maintain that it is unproblematic for conscious beings to overlap as long as they *share* their experiential properties with one another. If my head, for instance, is conscious, its consciousness makes no addition to mine: it and I share a stream of consciousness. This simply extends to conscious subjects a natural defence of manyism about non-mental objects.

3.1 Parsimony and Epiphenomenalism for Physical Objects

Lewis [1993] addresses the Problem of the Many by arguing that ‘the real opposite of identity is ... not distinctness in the sense of non-identity, but rather ... non-overlap’. When things overlap, they ‘are not entirely identical, not entirely distinct, but some of each’ [1993: 33]. And with larger or smaller overlap, things are more or less identical. This lets us interpret everyday talk of ‘two things’ as meaning not ‘two non-identical things’ but ‘two wholly or largely non-overlapping things’ (cf. López de Sa [2014]).

So ‘there is one chair here’ should be interpreted not as ‘there is a chair here which all chairs here are identical to’, but as ‘there is a chair here which all chairs here significantly overlap with.’ How much overlap is ‘significant’ is contextual, and it may be, as Sutton [2015] argues, that what matters is not the raw amount of overlap, but overlap in the parts serving crucial functions. For crowds and clouds, a bit of overlap is par for the course; for chairs, absolute disjointness is the common standard, and when there is overlap the key question is how many people can sit down. Either way, the apparent conflict between manyism and our everyday numerical judgements is resolved by replacing judgements of strict *distinctness* (non-identity) with judgements of different degrees of *disjointness* (non-overlap).

The spirit of this proposal is that our counts of objects primarily aim to organise, guide, and facilitate our navigation of the world. When I say ‘there’s a chair’, my primary aim is to direct people to some nearby matter and its suitability for sitting on. The difference between ‘two chairs’ and ‘one chair’ is primarily a difference in how many people can sit down. It matters little if, in metaphysical strictness, there

are really a trillion chairs, because those chairs are such that by sitting on one, you sit on them all.

Someone might object that manyism seems *unparsimonious*: Ockham's razor says not to multiply entities beyond necessity, and yet here we are, multiplying entities with wild abandon. They might also object that there's 'not enough work' for all these entities to be gainfully employed: anything one might cause, many others are equally suited to have caused, and so most will be epiphenomenal.

But this misunderstands composition. Things do not generally compete with their parts for causal efficacy, nor with overlapping things that share causally relevant parts. To use Merricks' example [2001: 111], the atoms in a baseball cause the shattering of a window, and commonsensically we think that the baseball causes exactly the same effect: this sort of 'overdetermination' is perfectly unobjectionable. (Merricks disagrees, and concludes that baseballs do not exist: this *ponens* is an excellent *tollens*.)

Similarly, wholes do not add to their parts when reckoning parsimony: it is not theoretical extravagance to suppose that, as well as bonded atoms, there also exist molecules, because the molecules are composed of the bonded atoms. Given parts, the wholes come for free, and trillions of free things cost no more than a few [cf. Blackmon 2021: 41-42]. This way to think about parsimony is surely preferable to eliminating molecules, baseballs, and tables from our ontology, like Sweeney Todd wielding Ockham's razor.

3.2 Can Experiences be Shared?

Suppose that manyists about physical objects can avoid epiphenomenalism and unparsimoniousness. Overlapping objects, in general, neither count for double nor causally compete. A defence of manyism about conscious subjects must extend this solution: overlapping subjects also neither count for double nor causally compete. For example, consider me and a subject composed of all of my atoms plus one more. In reckoning parsimony, they add only one atom to me; and our causing the same events is not overdetermination, because my causal powers are the same powers that they have.

Crucial to this claim is that the properties of the overlapping wholes are not multiplied: wholes share the very same property instances as their parts. In Sutton's [2014] terminology, the properties of overlapping things are 'non-summative': the amount they have together is the sum of their individual amounts minus whatever they share.

If your body weighs 150 lbs. and your arm weighs 9 lbs., then when you step on the scale, it registers 150 lbs. The scale did not forget your arm! The 9 lbs. of the arm is the same as 9 lbs. of the body. Weight is thus non-summative across non-identical objects where those objects share a supervenience base for weight [Sutton 2014: 622]

Sutton extends this principle to mental states, arguing that various Too-Many-Minds problems can be solved once we 'recognize that there are two non-identical beings that think the same thoughts' [2014: 622].³

³ This claim about thoughts raises semantic and epistemological challenges distinct from the phenomenological challenges I discuss: see Sutton [2014: 632-36] (cf. Unger [1980: 461-62], Merricks [2001: 103], Olson [2003: 329-31], and Weatherson [2003]). For the record, I consider Sutton's responses successful, and have little to add.

But can *consciousness* be non-summative? That is, can overlapping subjects share their token experiences? By ‘experience’ I mean a particular instance of a phenomenal property, a property defined by what it is like to instantiate it. We usually assume that each experience belongs to only one subject, because that seems central to the ‘privacy’ of experience, the fact that we know our own experiences in a different way than we know those of others. And this special epistemology seems essential to consciousness: it is how things seem *to me*, and thus tied to my perspective in a way that contrasts with the ‘public’, ‘objective’, nature of physical things.

If experiences, unlike instances of physical properties, cannot be shared, then manyism about minds will look very different from manyism about physical objects. Each of the trillion minds will have a stream of consciousness additional to that of the others. And a multiplicity of minds with their own, non-overlapping, sets of experiences, really would offend against parsimony, and really would struggle to assign each set of experiences non-redundant causal work to do.

I think this is why many philosophers have thought that manyism about minds is problematic, while manyism about physical objects is not: the objects can share their properties, but the minds cannot share their experiences. As Zimmerman puts it:

If there are two or three thinkers here, then there are two or three pains that have to be taken into account in utilitarian calculations, for instance; and two thinkers thinking ‘I am in pain’ ... [this] seems to introduce two or more pains, feelings of melancholy, etc., one for each coincident entity, where intuitively there should be only one. [Zimmerman 2005: 497]

I prefer to reason in reverse, so to speak: if minds are naturally occurring parts of the physical world, and if all parts of the physical world are susceptible to the division and overlap that generates multiplicity, we should see how far we can make sense of minds sharing their experiences just like objects share their physical properties.

Experience-sharing would revise our understanding of experiential ‘privacy’. We intuitively recognise an asymmetry between my access to my own experiences and others’ access to them—I know them from the inside, but others can only know them through their outside effects. But if experiences are shared, then ‘[two subjects] have privileged access to one another’s thoughts in virtue of those thoughts being numerically identical’ [Sutton 2014: 635].

But this need not erase the intuitive asymmetry. It remains true and significant that experiences are known in two different ways—‘directly’ and ‘indirectly’, we might say—and that a subject can only know their own experiences directly. Moreover, experiences are likely only shareable between mereologically overlapping subjects (as seems to be true of token property instances generally), and so my experiences are still directly knowable only to a very select group—me and the subjects who overlap with me. The objective facts of physical science, by contrast, are equally knowable by anyone, whatever their parts. So while we must revise our understanding of privacy, we need not give it up.

There are more subtle metaphysical objections to experience-sharing, which I have discussed at length elsewhere [Roelofs 2016: 3206–14, 2019: 57ff; Goff and Roelofs *forthcoming*].⁴ Since I cannot treat the topic thoroughly here, my defence of manyism is conditional: *if* two subjects can share experiences, *then* there is nothing implausible about manyism.

⁴ See also Basile [2010: 108ff], Coleman [2014: 34ff], Mørch [2014: 172–75], Blackmon [2021: 49].

4 Phenomenological Arguments Against Manyism

Suppose that experience-sharing is possible in principle. It might still be that something about our phenomenology tells us decisively that we are not sharing *our* experiences with anyone.

Phenomenological concerns seem to lie behind some support for Anti-Nesting principles (cf. the quotations collected by Blackmon [2021: 42-44]). For example, Tononi motivates his Exclusion Postulate as follows:

No matter how hard I try, I cannot become conscious of what is going on within the modules in my brain that perform language parsing ... [and] while I can interact with other people, I have no access to their internal workings. Exclusion applies also within consciousness: at any given time, there is only one consciousness ... – me – having one full experience, not a multitude of partial consciousnesses, each experiencing a subset of the contents of my experience. [Tononi 2012: 296]

I do not know exactly what Tononi's argument is here: why does denying the Exclusion Postulate entail access to our language parsing modules, let alone to other people's minds? But he does seem to suggest that if other subjects were sharing my experiences, then this fact should be somehow manifest: I should not be unaware of it 'no matter how hard I try' (I have encountered similar ideas in conversations about manyism). I disagree: a part of me might be undergoing some subset of my experiences without this being introspectively evidence to me. What can be said for this thought that manyism, were it true, should be introspectively evident?

4.1 The Subject-Revelation Argument

Many think that not only can a subject having an experience know that they are having it, but that they can thereby know the nature of that experience: undergoing pain teaches me exactly what it is for something to feel painful. This thesis is often known as 'revelation' [Johnston 1992: 223; Chalmers 2003; Stoljar 2006; Strawson 2015; Goff 2017: 109–10].

This thesis is often thought to pose problems for views according to which our experiences are different from how they initially seem, such as physicalism (on which experiences are brain processes or functional states) or constitutive panpsychism (on which experiences are assemblages of trillions of microexperiences). But we could formulate an analogous argument against manyism, as follows:

- (4) Having an experience puts me in a position to know its nature. (The Revelation Thesis)
- (5) For any subject *S* undergoing an experience *e*, being undergone by *S* is part of *e*'s nature.
- (6) *Therefore* In undergoing an experience, I am in a position to know of all the subjects undergoing it. (From 4 and 5)
- (7) Reflection on my experiences reveals no subjects undergoing them except me.
- (8) *Therefore* No subject except me is undergoing my experiences. (From 6 and 7)

This argument rests crucially on (5), that the subject(s) of an experience are part of its nature. Here is one reason to reject this premise: which subjects are undergoing a given experience may depend on facts external to that experience itself, or at least outside its

supervenience base. If both my whole body and my-body-minus-one-hair are sharing a certain brain-based pain, then the fact that these two distinct subjects are undergoing that pain depends on the existence of the hair: without it, these two subjects would not be distinct. But the hair is outside the supervenience base of the pain, so intuitively the pain's essential nature should not depend on it.⁵

Moreover, (5) seems unmotivated. Why think that all the subjects undergoing a certain pain are involved in that pain's essential nature? One might try to motivate it by stressing how intimately dependent experiences are on their subjects: experiences are simply modifications of subjects. This is the view Noonan expresses when he writes:

The concept of *someone's having* a perception is logically prior to the concept of *a perception* ... the relation between the self and its perceptions is analogous to that between the sea and its waves. The waves are modifications of the sea and perceptions are modifications of the self. [Noonan, 2003: 71]

The critic of manyism might point out that the essential nature of a wave *does* include its dependence on a medium: we could not understand what it is to be a wave without understanding that it is a wave *in* something. Perhaps we could not understand what *this particular* wave is, without understanding that it is a wave in *this particular* bit of the sea.

However, even this will not vindicate the argument against manyism. For a given wave on, say, a Croatian beach is at one and the same time a wave in the Adriatic Sea, and a wave in the Mediterranean Sea, and (if manyism is true) a wave in many other bodies of water we lack names for ('the Adriatic-and-Ionian-Seas-together'). But this roster of entities is *not* knowable just from scrutinising that wave, or its essential nature. That wave, plausibly, could have still existed if the rest of the Mediterranean had, a moment before, been deleted from existence (and the colossal shockwaves had not yet reached into the Adriatic), so that the wave was not in fact in the Mediterranean. All that its nature tells us about is its minimal supervenience base—the portion of water which minimally sufficient for that wave to exist. And knowing that minimal supervenience base does not tell us which larger entities might contain it.

4.2 The Special Unity Argument

A second phenomenological argument against manyism focuses on the overall *structure* of experience. This argument requires distinguishing between *total* and *partial* experience-sharing: between cases where two subjects have exactly the same set of experiences at a time, and cases where some experiences are shared, but one or both subjects also have some unshared experiences. Of the four sorts of Too-Many-Minds problems I went over in section 2, types 1 and 3 involve total sharing: if I and my brain, or I and me-minus-a-dubious-atom, share experiences, then plausibly we share all of our experiences. By contrast, problems of types 2 and 4 may involve only partial sharing: for example, if the left half of me, or my left hemisphere, are conscious subjects, then they presumably only undergo some, not all, of my experiences (although identifying which ones might be exceedingly complicated). Likewise, if

⁵ This claim, that an experience's subjects are not intrinsic to that experience, does not conflict with the idea that a subject's consciousness is intrinsic to that subject: it is consistent to hold that my experiences are 'in' me, but I am not 'in' my individual experiences.

some subsystem of my brain, or even each neuron, were itself conscious, then presumably its experiences would be some small subset of mine—although again, exactly how to map neural divisions onto experiential divisions might be very complex [Roelofs 2019: 166–68].

Mørch has argued specifically against *partial* experience-sharing, on the basis that if there were other minds partly overlapping with ours, that would be phenomenologically manifest in the structure of our experience:

... if there were overlapping experiences, then parts of my experience should be unified in a special way, or have a noticeable border around it that separates it from other parts of my experience ... [S]ubjects are not distinct from their experiences ... (and so) another subject of a part of my experience should show up in that part. (Hedda Hassel Mørch, personal correspondence, 2021)

I take the structure of this argument to be as follows.

- (9) If, within a set of experiences, some of them belong to a given subject and others do not, then those experiences should be unified in a special way.
- (10) No subset of my experiences is unified in a special way.
- (11) *Therefore* No subset of my experience belongs to another subject. (From 9 and 10)

One way to respond to this argument is denying (9), that experiences belonging to a single subject must always be ‘marked out’ by some special sort of unity. We might insist that the only kind of unity required for some experiences to belong to a single subject is phenomenal unity, mere being-experienced-together, and since all of my experiences are phenomenally unified, any subset of them is also phenomenally unified. For some or all of those subsets to be shared with another subject does not require them to be any *more* unified than that.

Alternatively, manyists might deny (10), the premise that human experience does not in fact show ‘special sorts of unity’. After all, human experience *is* intricately structured, into various sorts of clusters and groupings: different modalities, distinct aspects of our personalities, distinct attentional streams, distinct perceived objects, etc. The decision between these two approaches will depend on how we flesh out the special sort of unity that is required, and what it would look like, phenomenologically, for a subset of our experiences to display that unity. More demanding interpretations of ‘unity’ would make it easier to deny (9), while less demanding ones would make it easier to deny (10). A manyist could thus endorse a flexible combination of both responses: affirming that many subsets of our experiences do exhibit certain sorts of special unity, while denying that the experiences of an overlapping subject must exhibit any stronger sort of unity than this. Without saying more, I do not feel that a compelling case can be made that both (9) and (10) are true simultaneously, and thus manyism cannot be ruled out by an argument such as this.

5 Moral Arguments Against Manyism

Finally consider moral objections to manyism, specifically Simon’s ‘Hard Problem of the Many’ [2017]. Although manyism does raise difficult moral questions, they

divide into questions of everyday morality which manyists have good resources to address, and harder questions which are equally challenging for non-manyists.

5.1 Equal Treatment and Hedonic Beneficence

Simon first notes that, given natural variation in body size, there may be a different absolute number of conscious subjects present in the cloud we would have called ‘one person’ and the cloud we would have called another. This, he claims, forces manyists to choose between the following two theses.

Equal Treatment

‘Fully distinct (person-shaped) experiencer-clouds deserve equal moral treatment, other things equal.’ [Simon 2017: 452]

Hedonic Beneficence

‘If option φ relieves the pain of n experiencers, option ψ equally relieves the pain of m experiencers, and $n < m$, then other things equal one ought to do ψ rather than φ .’ [Simon 2017: 452]

Wherever a larger and a smaller person (as we would normally say) are both in pain, but resources to help them are scarce, Hedonic Beneficence seems to imply, if manyism is true, that we should prioritise relieving the pain of ‘the larger person’, since they are a cloud of (say) five trillion trillion experiencers, while ‘the smaller person’ is a cloud of only four trillion trillion experiencers. Yet this seems to violate Equal Treatment.

The obvious manyist response is that Hedonic Beneficence, and other everyday distributive principles, are correct, but only counting by disjointness rather than distinctness (that is, by almost-identity rather than identity). We should prefer to relieve the suffering of ‘more’ rather than ‘fewer’ non-overlapping subjects—that is, the suffering sets of conscious systems within which more conscious systems (with person-like minds) exist who do not overlap with one another. For clarity, we might distinguish two interpretations of Hedonic Beneficence.

HB-Distinct

If option φ relieves the pain of n distinct experiencers, option ψ equally relieves the pain of m distinct experiencers, and $n < m$, then other things equal one ought to do ψ rather than φ .

HB-Disjoint

If option φ relieves the pain of n disjoint (non-overlapping) experiencers, option ψ equally relieves the pain of m disjoint experiencers, and $n < m$, then other things equal one ought to do ψ rather than φ .

Interpreting Hedonic Beneficence as expressing HB-Disjoint is not an objectionable ‘revision’ to everyday morality: everyday morality emerges from the everyday experience of engaging with others, and whatever default authority it has comes from that experience. If manyism is true, then that experience has always been experience of engaging with clouds of overlapping subjects, not with single subjects. HB-Disjoint is then the natural theoretical interpretation of the intuitive principles that emerge from that experience. HB-Distinct only gets to be the default interpretation of everyday morality if our everyday experience has been with strictly single subjects. To reject

manyism because it forces a ‘revision’ to Hedonic Beneficence is thus question-begging.⁶

Simon foresees this response, but claims it

has troubling implications If experiencers matter less when they almost entirely overlap other experiencers, do they also matter less when they only overlap other experiencers a little bit? Where do we draw the line? Do [conjoined] twins matter less before they are successfully separated than after? [Simon 2017: 454]

But this is a misunderstanding: the claim is not that overlapping another subject somehow ‘subtracts from’ someone’s moral status. Inferring that subjects must ‘count for less’ when they overlap requires assuming that since they collectively count for one, each must be counting for some fraction of one, because their status is being ‘added up’. But the whole thrust of my defence of manyism has been that they combine non-additively. They each have full moral status, but they share this moral status with their overlappers.

Moreover, HB-Disjoint is just a formalised version of an everyday moral rule of thumb, for dealing with different numbers of wholly disjoint subjects (separate person-clouds). It is not, and I will not try to provide, a systematic explanation of the basis for the equal moral status of separate person-clouds, or how different sorts of overlap interact with that basis. But one obvious thought is that overlap matters most in the brain, the supervenience base for complex human-like consciousness. If two conjoined twins share legs but not a brain, for instance, manyists can trivially derive the obvious right answer: they count as two people, not as one.⁷ More broadly: we would have to look not just at bare distinctness or disjointness, but at the bases for morally relevant capacities (this is effectively what Sutton [2015] calls ‘counting by functional almost-identity’). If what makes a subject matter morally is their capacity to suffer, or reason, or form relationships, then what matters for counting is overlap or non-overlap in the brain systems responsible for suffering, or reasoning, or relating.⁸

5.2 The Morality of Overlap

Simon suggests a more hypothetical worry: what about two subjects whose brains are neither disjoint (or nearly-disjoint) nor identical (or almost-identical), but somewhere in between? Some hypothetical forms of craniopagus conjoinment might have this form, if the twins shared, so to speak, ‘half their brains’, and if conscious artificial intelligences are created, they might be formed into, or form themselves into, this sort of configuration. To put the mereological issues precisely, the concern is with cases

⁶ Similar remarks apply to Johnston’s personite problem: if personites exist, our moral intuitions grew from engagement with sequences of overlapping, psychologically continuous, person-like beings, and should be formulated appropriately. See Johnston [2017b: 211–14] for a discussion of this type of response, which Johnston claims cannot save commonsense morality, and is still ‘unfairly biased against the personites’ [2017b: 214]. I am inclined to disagree, but this raises complex questions about persistence and survival across time that go beyond what I can address here.

⁷ As best I know, this applies to all actual conjoined twins: even where two brains are linked by nerve tissue, they are either not overlapping, or overlap only slightly. For discussion of the issues raised by Krista and Tatiana Hogan, the most well-studied pair of craniopagus conjoined twins, see Dominus [2011], Langland-Hassan [2015], Roelofs [2019: 11, 112 fn. 57], Cochrane [2021].

⁸ Simon [2017: 454] notes that there are also trillion experiences wherever we thought there was one. This is fine: experiences too can be counted by overlap, almost-identity, or functional almost-identity.

where the manyist would say (i) there is no conscious intelligent subject with whom all subjects present overlap almost completely, but (ii) there are conscious intelligent subjects with whom all subjects present overlap substantially. Call this a ‘Y-shaped’ situation (the subjects are like two lines partly overlapping to form a ‘Y’). The problem of how to make sense of the seemingly ‘indeterminate or intermediate’ number of subjects present, is what Schwitzgebel [2021] calls the ‘continuity problem’.

I don’t think we can or should resolve the continuity problem in the abstract: should Y-shaped situations arise, it will be a genuine challenge to work out how they fit with our intuitive moral systems. My remarks above suggest initial guidelines: to pay attention to capacities like suffering or autonomous choice and see whether there is overlap of the neural systems responsible for these capacities. But we understand so little of the neurology of these capacities, even in typical human brains, that this is a suggestion, not a solution.⁹ Moreover, the unusual structure of these systems may itself change the neurology of these morally capacities: the architecture of autonomous choice may develop in plastic ways in the presence of another such architecture, so even if the subjects start out non-overlapping in a particular regard, they may come to overlap in the process of living together.

How to relate to Y-shaped systems will be a moral challenge for us and also for them, since one crucial issue is how they think of themselves: do they see themselves as a single system with a unified viewpoint and goals, or as two intimate companions, sharing a bond beyond anything we can approximate with our words for romance? Or do they hate each other, and see their overlap as a curse? Or do they relate in a way that demands the fashioning of its own conceptual vocabulary? Our primary moral directive, surely, is to ask the question and try to respect the answers we get, not to devise answers of our own on theoretical grounds.

The likeliest candidate for an actual Y-shaped case is the split-brain phenomenon, which some theorists argue should be seen as two distinct subjects, whose brains each consist of one cerebral hemisphere and the shared subcortical brain structures [Schechter 2014, 2015, 2018]. If that is correct, there are two subjects here who overlap very significantly but nowhere near completely, and no single subject that overlaps them both near-completely. Of course, the empirical data about these cases is mixed and hard to interpret [Pinto *et al.* 2017], but it may provide a useful reference point.

As far as I know nobody suggests that split-brain patients should be treated, in any moral sense, as two people. Schechter herself insists that despite being two minds, they are collectively only one ‘interpersonal agent’ [2018: 183ff]. But there are specific features of this case that matter: for a start, the capacity for choice seems to be shared, in that any choice made by one subject is *ipso facto* made by the other (at least, conflicts seem to be rare). It is also possible that the neural basis for some forms of somatosensory suffering is shared. Moreover, crucially, even if there are two people here, they don’t think of themselves as two, but as one [Schechter 2018: 156–80]. The split-brain patient is remarkable because the partial dissociation between their minds requires special techniques to demonstrate it, and is not recognised introspectively. They act like, talk like, and see themselves as, a single person. We should defer to this self-perception in the absence of a strong reason not to.

⁹ For some discussion of the ethics of partly-unified, partly-disjoint systems somewhat similar to this, see Churchland [1981: 87–88], Rovane [1998: 141], Hirstein [2008], and Roelofs [2019: 270–94].

But this case is deeply ambiguous: quite possibly there *is* a single encompassing subject here, despite its psychological dissociations. Even if this case is genuinely Y-shaped, its important features might differ from other cases, offering little prospect of a general rule.

It is not a shortcoming of a theory of mind that it does not answer these questions for us in advance, but simply alerts us to how difficult they are likely to be. And non-manyist approaches, like Simon's own dualism, are no better off: rather than grappling with the moral complexity that such cases would present, they demand *a priori* that, since such complexity is unattractive, there must always be a non-physical fact of the matter about the precise number and boundaries of subjects, even if we cannot detect that fact empirically. If we prefer to face what is conceptually difficult but empirically tractable over positing a clean answer that can never be discovered, we should be open to manyism.

6 Conclusion

It would be surprising if wherever we thought there was a physical object, there were a trillion almost-exactly overlapping objects. It would be even more surprising if this turned out to be true of minds! But if overlapping objects neither compete for causal efficacy nor add up for parsimony calculations, we have no reason to think this surprising claim is absurd. I argue this even goes for minds: no number is too many, if they overlap sufficiently. Metaphysical, phenomenological, and moral objections all fail: properly understood, manyism about minds is compatible with our best available metaphysics, with our everyday moral practice, and with the phenomenology of self-awareness. This does not establish the truth of manyism: it just shows that if an independently attractive theory entails manyism, that is no reason to reject it.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

References

- 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.
- 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
- Chalmers, David 2003. The Content and Epistemology of Phenomenal Belief, in *Consciousness: New Philosophical Perspectives*, eds Q. Smith and A. Jokic, Oxford: Oxford University Press: 220–272.
- Churchland, Paul 1981. Eliminative Materialism and the Propositional Attitudes, *Journal of Philosophy* 78/2: 67–90.
- Cochrane, Tom 2021. A Case of Shared Consciousness, *Synthese* 199/1–2: 1019–37.
- Coleman, Sam 2014. The Real Combination Problem: Panpsychism, Micro-Subjects, and Emergence, *Erkenntnis* 79/1: 19–44.
- Dominus, Susan 2011. Could Conjoined Twins Share a Mind? *New York Times Magazine*, May 29.
- Goff, Philip 2017. *Consciousness and Fundamental Reality*, New York: Oxford University Press.
- Goff, Philip and Luke Roelofs forthcoming. In Defence of Phenomenal Sharing, in *The Phenomenology of Self-Awareness and Conscious Subjects*, eds Julien Bugnon, Martine Nida-Rümelin, and Donnchadh O'Conaill, London: Routledge.

- Hirstein, William 2008. Mindmelding: Connected Brains and the Problem of Consciousness, *Mens Sana Monographs*, 6/1: 110–30.
- Johnston, Mark 1992. How to speak of the colours, *Philosophical Studies* 68/3: 221–63.
- Johnston, Mark 2017a. The Personite Problem: Should Practical Reason be Tabled?, *Noûs* 51/3: 617–44.
- Johnston, Mark 2017b. Personites, Maximality, and Ontological Trash, *Philosophical Perspectives* 30/1: 198–228.
- Johnston, Mark 2021. The Subject and its Apparatus: are they Ontological Trash? *Philosophical Studies* 178/8: 2731–44.
- Kammerer, François 2015. How a Materialist Can Deny That the United States is Probably Conscious—Response to Schwitzgebel, *Philosophia* 43/4: 1047–57.
- Langland-Hassan, Peter 2015. Introspective Misidentification, *Philosophical Studies* 172/7: 1737–58.
- Lewis, David 1976. Survival and Identity, in *The Identities of Persons*, ed. Amélie Oksenberg Rorty, Berkeley, CA: University of California Press: 17–40.
- Lewis, David 1993. Many, but Almost One, in *Ontology, Causality, and Mind: Essays in Honour of D.M. Armstrong*, eds J. Bacon, K. Campbell, and L. Reinhardt, Cambridge: Cambridge University Press: 23–45.
- López de Sa, Dan 2014. Lewis vs Lewis on the Problem of the Many, *Synthese* 191/6: 1105–17.
- Merricks, Trenton 2001. *Objects and Persons*, New York: Oxford University Press.
- Mørch, Hedda Hassel 2014. *Panpsychism and Causation: a New Argument and a Solution to the Combination Problem*, PhD thesis, University of Oslo.
- Mørch, Hedda Hassel 2019. Is Consciousness Intrinsic? A Problem for the Integrated Information Theory, *Journal of Consciousness Studies* 26/1–2: 133–62.
- Noonan, Harold 2003. *Personal Identity*, 2nd ed., London: Routledge.
- Olson, Eric 2003. An Argument for Animalism, in *Personal Identity*, eds R. Martin and J. Barresi, Malden, MA: Blackwell: 318–35.
- Paul, L. A. 2002. Logical Parts, *Noûs* 36/4: 578–96.
- Paul, L. A. 2006. Coincidence as Overlap, *Noûs* 40/4: 623–59.
- Pinto, Y., D. Neville, M. Otten, P. Corballis, V. Lamme, E. de Haan, N. Foschi, and M. Fabri 2017. Split Brain: Divided Perception but Undivided Consciousness, *Brain: A Journal of Neurology* 140/5: 1231–37.
- Putnam, Hilary 1965/2003. The Nature of Mental States, in *Philosophy of Mind: Contemporary Readings*, eds T. O'Connor, D. Robb, and J. Heil, London: Routledge: 210–21.
- Roelofs, Luke 2016. The Unity of Consciousness, Within and Between Subjects, *Philosophical Studies* 173/12: 3199–221.
- Roelofs, Luke 2019. *Combining Minds: How to Think About Composite Subjectivity*, New York: Oxford University Press.
- Rovane, Carol 1998. *The Bounds of Agency: An Essay in Revisionary Metaphysics*, Princeton, NJ: Princeton University Press.
- Schechter, Elizabeth 2014. Partial Unity of Consciousness: A Preliminary Defense, in *Sensory integration and the unity of consciousness*, eds D. Bennett, D. J. Bennett, and C. Hill, Cambridge, MA: MIT Press: 347–74.
- Schechter, Elizabeth 2015. The Subject in Neuropsychology: Individuating Minds in the Split-Brain Case, *Mind and Language* 30/5: 501–25.
- Schechter, Elizabeth 2018. *Self-Consciousness and 'Split' Brains: The Minds' I*, Oxford: Oxford University Press.
- Schwitzgebel, Eric 2015. If Materialism Is True, the United States Is Probably Conscious. *Philosophical Studies* 172/7: 1697–1721.
- Schwitzgebel, Eric 2021. Review of *Combining Minds: How to Think About Composite Subjectivity*, *Notre Dame Philosophy Reviews*, <https://ndpr.nd.edu/reviews/combining-minds-how-to-think-about-composite-subjectivity/>
- Seager, William 2010. Panpsychism, Aggregation and Combinatorial Infusion, *Mind and Matter* 8/2: 167–184.
- Seager, William 2017. Panpsychist Infusion, in *Panpsychism: Contemporary Perspectives*, eds G. Brüntrup and L. Jaskolla, Oxford: Oxford University Press: 229–48.
- Sider, Theodore 2003. Maximality and Microphysical Supervenience, *Philosophy and Phenomenological Research* 66/1: 139–49.

- Simon, Jonathan 2017. The Hard Problem of the Many, *Philosophical Perspectives* 31/1: 449–68.
- Stoljar, Daniel 2006. *Ignorance and Imagination*, Oxford: Oxford University Press.
- Strawson, Galen 2015. Self-Intimation, *Phenomenology and the Cognitive Sciences* 14/1: 1–31.
- Sutton, C.S. 2014. The Supervenience Solution to the Too-Many-Thinkers Problem, *Philosophical Quarterly* 64/257: 619–39.
- Sutton, C.S. 2015. Almost One, Overlap and Function, *Analysis* 75/1: 45–52.
- Tononi, Giulio 2012. Integrated Information Theory of Consciousness: An Updated Account, *Archives Italiennes de Biologie* 150/2–3: 56–90.
- Tononi Giulio, and Christoph Koch 2015. Consciousness: here, there and everywhere? *Philosophical Transactions of the Royal Society B*, 370/20140167: 1–18.
- Unger, Peter 1979. Why There are No People, *Midwest Studies In Philosophy* 4/1: 177–222.
- Unger, Peter 1980. The Problem of the Many, *Midwest Studies in Philosophy* 5/1: 411–68.
- Unger, Peter 2004. Mental Problems of the Many, in *Oxford Studies in Metaphysics, Volume 1*, ed. Dean W. Zimmerman, Oxford: Clarendon Press: 195–222.
- Weatherson, Brian 2003. Many Many Problems, *Philosophical Quarterly*, 53/213: 481–501.
- Zimmerman, Dean W. 2005. Material People, in *Oxford Handbook of Metaphysics*, eds Michael J. Loux and Dean W. Zimmerman, Oxford: Oxford University Press: 491–526.
- Zimmerman, Dean W. 2010. From Property Dualism to Substance Dualism, *Aristotelian Society Supplementary Volume* 84/1: 119–50.