Religious Studies (2025), 1–19
doi:10.1017/S0034412525100942

CAMBRIDGE UNIVERSITY PRESS

ORIGINAL ARTICLE

Does panpsychism entail anti-realism? The worm in the panpsychist apple

Miri Albahari

Department of Philosophy, The University of Western Australia, Crawley, WA, Australia
Email: miri.albahari@uwa.edu.au

(Received 10 July 2025; revised 10 July 2025; accepted 10 July 2025)

Abstract

Panpsychists commonly hang onto the 'realist' assumption that our world with its structures has an observer-independent, often spatial element to it, even while they claim that those structures are realized by the experiences of subjects. I argue that this assumption is the 'worm in the apple' that lurks behind two of panpsychism's major problems: the subject (de)combination problem and what I call the 'inner-outer gap problem'. Abandoning this assumption sidesteps those problems, but commits panpsychism to anti-realism.

Keywords: panpsychism; (de)combination problem; inner-outer gap problem; realism; anti-realism

Introduction

Panpsychism is the view that consciousness is fundamental to physical reality. Modern versions add to it the idea that our conscious experience is derived from experiences had by subjects that comprise the intrinsic natures of the most basic physical elements, whether on a micro or cosmic level. Panpsychism is increasingly seen as a viable avenue along which to dodge the bullets that target its well-known rivals: dualism and materialism. But it suffers from what is known as the combination (and decombination) problem.[1] Its most pressing version concerns the idea of subjects sharing their conscious contents with each other. While a subject is a notoriously mysterious entity, one of its defining features seems to be the utter privacy of its conscious experiences. I might share a chocolate bar, but no other subject gets to share the very taste that I enjoy. As William James (1890) observes, the boundaries between subjects seem the most unbreachable to be found in nature. Yet the very nature of panpsychism seems to demand that the boundaries of subjects(s) be breached, their experiences shared or passed on to the subjects whose existence they comprise.

It is not only the combination problem that threatens panpsychism. I have argued that what I have called the 'inner-outer gap problem' is every bit as devastating (Albahari 2022). While matter and its structures are supposed to be realized by conscious subjects, many panpsychists still write as if the structures existed unproblematically out there, in space, as observer-independent, *sub species aeternitatis* features of the world. For instance, if Jim examines what he takes to be Lucy's brain, its structures are presumed to exist, prior to

his examination of them, in a region of space literally occupied by Lucy's head – whilst somehow also being implemented by Lucy's private experiential states (and perhaps those of her neurons). But then how can the inner experiences enjoyed by Lucy (or her neurons) become without mystery a raft of locational observer-independent structures and behaviours that exist out there in space? The gap that exists between mind and matter yawns as wide as the gap that creates problems for materialism and dualism. Instead of mind emerging mysteriously from matter, or causally setting into motion its physical behaviours, panpsychism now has matter and its structures emerging mysteriously from mind, setting into motion its physical behaviours. The mind-matter gap inherent in this version of panpsychism is symmetrical to that suffered by materialism and dualism, undermining a primary motivation to be a panpsychist.

Could these two problems be connected? My 2022 article closed by suggesting that the pervasive assumption which lurks behind the inner-outer gap problem could also lurk behind the combination problem. The assumption, likely inherited from a physicalist or dualist paradigm, is that physical reality involves structures, compositional relations, and behaviours that exist in an outer *observer-independent* manner. Although the structures are supposed to be realized by minds, they are frequently talked about as if they existed in a straightforwardly external, typically spatio-temporal manner – in line with how the world seems. This assumption is sometimes aligned with a stance that is described as *realist* as opposed to *anti-realist* (Chalmers 2020). The current article makes a proper case for the claim that this assumption of observer-independence is the worm in the panpsychist apple, one that underpins both the (de)combination and inner-outer gap problems. I suggest that once we expel the worm, we are left with an anti-realist and more classically idealist version of panpsychism that does not suffer from any of the mind-matter gaps which create problems for materialism, dualism, and standard panpsychism. But the position does force us to rethink subjects in relation to the world. Subjects cannot, for instance, be literally embodied creatures in a spatio-temporal universe with actual mind-independent locations. Subjects will still derive their experiences from other subjects, but not in the way panpsychists have usually conceived. While this is a view I defend elsewhere, especially in developing the metaphysics of Advaita Vedanta (Albahari 2019, 2024), it is not the goal of this article to show that anti-realist panpsychism is, all things considered, a stronger position. The primary goal is rather to point out that avoiding the (de)combination and inner-outer gap problems commits panpsychism to anti-realism.

The article will proceed as follows. I'll start by briefly outlining the problems for materialism and dualism, before going through what is known as the Russellian intrinsic nature argument in order to show how panpsychism offers a resolution. Next, I'll introduce the subject combination, decombination, and inner-outer gap problems for panpsychism, with a view to arguing that the underlying assumption of an observer-independent world – leant on heavily in the intrinsic nature argument – underpins all the problems. Then, I'll sketch the outline for a version of anti-realist panpsychism. Following this, I'll address an objection that is motivated by the work of Galen Strawson and David Chalmers, according to which realist panpsychism does not have to succumb to the inner-outer gap problem. According to their objection, the ideas of physical, spatially located matter can be understood as a version of structural realism that allows the external world and its structures to be grounded in the minds of distal non-human observers, while being observer-independent in relation to humans. I'll argue that this approach merely sharpens the dilemma: panpsychism must either give up on the idea that scientific claims about the external world are literally true or fall prey to the inner-outer gap problem. I'll close by suggesting that realist panpsychism faces the same underlying metaphysical challenges as anti-realist versions, which may favour, all things being equal, an anti-realist approach.

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

Setting the scene: materialism, dualism and their objections

Mind and matter seem like very different phenomena. Mind involves a conscious subject which undergoes first-personal experiences that are replete with such qualities as the taste of apple, its enjoyment, and the desire for more. By contrast, matter is usually characterized in third-personal, mechanistic, structural terms, such as those describing patterns of electrochemical discharge in the brain. While there are some observable correlations between the two, there is a longstanding puzzle as to how mind and matter exactly relate to one another. This is known as the mind-body problem. The two main kinds of answer, over the past few centuries, have come from dualist or materialist camps.

Dualism is the view that the mind or its properties are fundamentally distinct from material properties, just as they appear to be. Dualists who want to preserve the intuitive idea that mental states cause physical behaviours are faced with what is known as the causal exclusion problem. It is a common scientific assumption that the physical universe is causally closed, with every physical event having a sufficient physical cause. If I lift up an apple to take a bite, then this behaviour, while seemingly propelled by a desire for apple-taste, is fully caused by neurophysiological events that run in tandem with the desire. The desire for an apple is squeezed out of its causal role, becoming epiphenomenal – a consequence many find unacceptable.[2]

Does materialism fare better? Here, mental properties are regarded as identical to or emerging from physical properties of the brain. The desire for an apple, turning out to be none other than the underlying brainstates or their emerging properties, can now cause my biting into it. But this raises a new problem. It seems conceivable that one could grasp in all its intricacies the physical story of how my brain-states lead to the biting of an apple and be none the wiser of the fact that they are accompanied by an appley taste. That a complete physical account of these brain-states would shed no light on *why* they are associated with experiential qualities seems to reveal an explanatory gap between matter and mind. This is known as the hard problem of consciousness (Chalmers 1996).

While each side has attempted to address the objections through more refined variants of their positions, the general consensus is that the problems haven't really evaporated. This has led philosophers to seek a resolution elsewhere by taking what they see is an entirely different approach to the mind-body problem. Some are returning to a line of argument that was set in motion by Bertrand Russell.

The intrinsic nature argument for panpsychism

In 1927, Russell made an observation about matter that would become a step in one of the most important contemporary arguments for panpsychism. The observation, made also by Arthur Eddington and Stephen Hawking, is that matter is characterized in physics not by what it intrinsically is but by what it does. Take electrons. What we may naively think of as tiny billiard balls are actually described in abstract, mathematical, and relational terms: how they are disposed to exert variously charged forces on other particles in accordance with lawlike equations. The intrinsic stuff that implements these behaviours is left out of the equation. So, what could this stuff be? Or to quote Hawking, 'What is it that breathes fire into the equations and makes a universe for them to describe?' (1988, 174). Some are happy to say that there is no such fire: the physical world is abstractions, equations, and dispositions all the way down. However, many philosophers, including Russell himself, are not content with this line of thinking, which amongst other things seems to conflate an intuitively weighty distinction between the concrete and the abstract.

If we accept that there has to be a categorical intrinsic nature that realizes the abstract equations, the question becomes: what is the best candidate for this nature? One proposal

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

casts it as a neutral medium that underpins the mind-matter correlations while being neither mental nor physical. The medium might be completely unknowable, or it might be a non-conscious but familiar quality that isn't exhaustively structural, like the colour and sound qualities that show up through our experience.[3] This however leaves a gap between the neutral medium and the conscious states. How, for instance, would the non-conscious qualities generate that peculiarly subjective character of our experiences by which the different qualities seem presented together in a unitive conscious field to a single point of view we think of as belonging to a subject?

There is another candidate. While we present to the outside world as a body and brain with identifiable structures and behaviours, we know that there is more to us than just this. There is something it is like to be us *from the inside* as we execute our behaviours. Putting these ideas together, panpsychists suggest that in our capacity as conscious subjects of experience, *we as subjects* have the kind of intrinsic categorical nature that can realize the externally observable behaviours and structures of our bodies and brains. Applying this idea more widely, what we encounter as a physically measurable world from the outside, whether simple or complex, is also none other than the activities of various subjects as they undergo experiences. Casting conscious subjects as the intrinsic nature of matter seems not only to avoid any gap between neutral and conscious entities; it is also parsimonious, not multiplying types of intrinsic nature beyond what is necessary. We already know that we have an inner conscious life that outstrips all the equations, so why not cast everything as having the same underlying nature?

Approaching panpsychism via the Russellian intrinsic nature argument seems to promise a way out of both the causal exclusion problem and the hard problem. While details will need finessing, a resolution to the hard problem seems there in the idea that all matter is consciousness-involving. Consciousness doesn't mysteriously come into being at some arbitrarily high level of complexity. And the causal exclusion problem seems ameliorated by allowing conscious states a central role in the causal nexus via its implementation of the physical equations. The reason physicists don't directly encounter this categorical nature in their findings is that the conscious subjectivity implementing the equations is by its very nature hidden from external purview. A neurophysiologist observing our brain and bodily behaviour from the outside would likewise encounter only the mechanical manifestations of what we know, from inside, to be a lot more than just this.

Panpsychism and its variants

Panpsychists are not committed to saying that every material entity is a subject, or that the inner conscious lives of non-biological entities resemble our own. Perhaps the experiences are extremely simple – primal pulses of attraction or repulsion – or they are exceedingly complex. There are theories in circulation as to which kind of processes facilitate subjecthood, with items such as quarks, atoms, neurons, or the cosmos more likely to feature than tables, rockets, or rocks. Assuming one of these theories to be viable, the question becomes: which subjects are the most basic? Independent scientific and philosophical considerations are brought into play here, with many opting for the historically popular line that fundaments coincide with the very small. A few, however, are persuaded by Jonathan Schaffer's (2010) arguments for priority monism, which claim the whole cosmos to be the most fundamental and unified concrete entity from which other physical particulars are derived. Panpsychists who think that physical reality bottoms out at the quantum level are known as micropsychists, while those who think that it tops out at the cosmic level are known as cosmopsychists.[4]

Having decided on where the fundaments lie, panpsychists now focus on the nature of the relation between the fundamental level of consciousness and that of our own

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

consciousness. Just as there is assumed to be some kind of physical-composition relation between the fundament(s) and middle-sized entities such as us, there is assumed to be some kind of experiential-composition relation between the subjective experiences had by the fundament(s) and our own. Constitutive panpsychists usually hold this experiential relation to be one of grounding. If the grounding is complete, the appropriately related experiences of the fundament(s) will combine or break down to guarantee metaphysically that we, as macro-subjects, have the experiences that we do, perhaps via some intermediate subject-entities such as neurons. Emergent panpsychists, by contrast, hold that laws of nature are additionally needed to connect our experiences with those of the fundament(s). Panpsychists also differ over the degree to which the physical world is consciousness-based, with 'pure' or 'idealist' variants of panpsychism basing physical reality entirely in conscious experiences, and 'impure' variants incorporating non-conscious elements such as spatial properties. My discussion will focus on the pure variants as these more obviously promise to close any mind-matter gaps.

Philosophers will debate which version best navigates both the hard problem and the causal exclusion problem. David Chalmers (2016) has argued that a Russellian variant of constitutive panpsychism works best, while Philip Goff (2024) argues that some hybrid of constitution and emergence is to be preferred. Regardless of their details, both versions, I'll soon argue, are infested by the worm in the apple. Neither closes the mind-matter gap that besets materialism and dualism. But to expose the worm I must first outline some major problems for panpsychism, starting with the most notorious, the subject combination and decombination problems.

The subject combination and decombination problems

Fully appreciating the subject combination and decombination problems requires prior reflection on what a subject is. Although subjects have long been a source of philosophical puzzlement, a subject, such as ourselves, is at the very least a centred conscious perspective to which observable objects are presented. While objects will include such things as apples and trees, the most immediate items to impinge upon our consciousness are an array of sensory and cognitive qualities that I have termed *cognisensory imagery* – objects such as the look, taste, or thought of an apple (Albahari 2019).

With this in mind, we can notice at least two striking facts about subjects that are behind their historically puzzling nature. First, a subject can never, in its capacity as a conscious, observing perspective, seem to become the focal object of its own or another subject's observation. (The puzzle: what is the nature of this observationally elusive subject?)[5] Second, while two subjects might observe the same external object, such as an apple, they can never appear to share the very same instance of cognisensory imagery through which the external object is observed. Such privacy of imagery to a subject's perspective serves indeed to individuate one subject from another. If I am aware of a desire for an apple, this cognitive imagery is not in the conscious field of any other perspective, and nor is any other perspective's imagery in my conscious field. This creates an unimpeachable sense of separation between myself – as owner of the imagery – and other subjects. On the privacy of a subject's thoughts and experiences in relation to its boundaries, William James has written:

> Each of these minds keeps its own thoughts to itself. There is no giving or bartering between them. No thought even comes into direct *sight* of a thought in another personal consciousness than its own. Absolute insulation, irreducible pluralism, is the law ... Neither contemporaneity, nor proximity in space, nor similarity of quality and content are able to fuse thoughts together which are sundered by this barrier of

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

> belonging to different personal minds. The breaches between such thoughts are the most absolute breaches in nature [James 1890, 226].

When panpsychists talk about the combining of experiences, whether through a process of merging, sharing, or passing on, the subject's perspectival architecture (as we might call it) is violated on at least one of these major fronts. Take a typical version of micropsychism, along with a supposition that our neurons are individually conscious. The combination of individual neuronic experiences into our own will demand that the boundaries of these microsubjects are breached and their contents spilled into ours, thereby compromising the privacy that separates one subject from another.

Cosmopsychism, while avoiding the combination problem, seems to do no better. The cosmic subject's global experience must yield our narrower perspectives with our own experiences. But a 'decombination' problem now looms, again forcing a compromise of the subject's perspectival architecture. Philip Goff (2017a) has for example defended a version of cosmopsychism in which the cosmic experience literally contains all of us having our experiences, making our perspectives objects of cosmic purview. He later modifies his position to deny the subsuming of perspectives only to go in for 'qualia transference' by which the cosmic subject's cognisensory imagery is literally passed onto us from the cosmos like a baton (2024). While Goff (with Luke Roelofs, Goff and Roelofs forthcoming) goes to some lengths to defend the coherency of phenomenal sharing, he has literally *changed the subject* such that its target is no longer recognizable as that observationally elusive and private entity which has puzzled generations of philosophers.[6] It is an exercise that also seems *ad hoc*, with its main purpose being to save panpsychism.[7]

There are of course other versions of panpsychism besides these. But so long as the subject's perspectival architecture is being compromised, there will remain combination and decombination problems. While many are content to live with this compromise, it would be better, if we could, to live without it.

A problematic assumption behind the (de)combination problem

I conjecture that the (subject) combination and decombination problems stem from how we – including panpsychists – implicitly picture the combination of experiences, even if denying that such a picture is literally true. The experiences and perspectives of 'smaller', for example neural, subjects are imagined to be in spatial proximity to one another, with their boundaries, like soap-bubble membranes, somehow popping and merging their contents into our own experience. Alternatively, if the cosmos is conscious, we envisage our perspectives and experiences as somehow contained inside of its perspective and experience. From the outset we are envisaging subjects as physically embodied experience-containing vessels that mix their contents together to make new experiences, the combining taking place inside the locations of the vessels.

We envisage experiential combination in this spatially oriented, boundary-breaching way, I surmise, because the physical entities, to which we ascribe such experiences, are assumed to be spatial objects whose parts can come together in such a manner. Whether we imagine neurons combining into salient networks, or atoms into molecules, or tributaries into a river, or Lego blocks into a model of the Eiffel Tower, the compositional relations between physical entities are generally taken to hold between non-mental particulars that occupy space in accordance with various laws of nature. The smaller elements combine into the larger, sometimes ceding their individual boundaries.

From these reflections, I thereby contend that the line of thought leading to the combination and decombination problems is anchored in a two-fold assumption that panpsychists and their critics tacitly adopt, *despite* panpsychism's putative grounding of

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

physical structures in the experiences of subjects. First, panpsychism continues to be conceived of as a position that places the world's physical objects and its structures in observer-independent regions of space – in just the same way a materialist or dualist would. And second, the experiences imputed to these physical entities are envisaged to combine in a way that aligns with the spatial limitations of their hosts. If Jim observes Lucy's brain as she enjoys an apple, for instance, it is typically assumed that the neural structures Jim observes are literally inside of Lucy's skull, combining into intelligible, law-governed patterns within regions of mind-independent space. And by picturing the experiences as annexed to the spatially proximate vessels of neurons and brains, we are invited to imagine the experiences as combining in a way that parallels their physical bearers. We can see these spatially oriented assumptions at work in the following passage about micropsychism from the Stanford Encyclopedia of Philosophy entry on Panpsychism:

> According to constitutive micropsychism, micro-level entities have their own very basic forms of conscious experience, and *in brains* these micro-level conscious entities somehow come together to constitute human and animal consciousness (Goff et al. 2022, my emphasis).

Just as the spatially proximate neurons are assumed to combine into intelligible physical networks, the experiences ascribed to these neural subjects are thought to combine into the intelligible subjective unity that we call ourselves – a line of thinking that walks right into the combination problem. A similar progression of thought often runs through cosmopsychism. By picturing ourselves as physical entities located inside a spatially extended cosmos, it is erroneously supposed that our conscious experience must also be contained inside the cosmic conscious experience. This walks right into the decombination problem.

Is there a way of envisaging constitutive relations between experiences that doesn't get ensnared by the (de)combination problems? Let us first see how this very same assumption of an observer-independent physical world plays into what I've called the inner-outer gap problem.

The inner-outer gap problem

Return to Jim's observation of Lucy eating an apple. As panpsychism has so far been characterized, Lucy's private subjective states, such as her desire for the next bite, are what supposedly animate her external apple-crunching behaviour. And what realizes the neurological structure that Jim observes, as he peers into Lucy's brain, is supposedly Lucy's experience, perhaps in virtue of those belonging to simpler entities such as her neurons. When this characterization combines with the entrenched presupposition that continues to regard brains, physical structures, and behaviours as if they were straightforwardly spatial and observer-independent entities, the inner-outer gap problem arises. It is simply this: *How could inner subjective experiences also, without mystery, be, become, or elicit an edifice of outer, observer-independent locational structures and behaviours?* This manifestation of outer, observer-independent, spatial structures and behaviours from inner, first-person subjective experiences is no less mysterious and gappy than the manifestation of subjective experiences from material brain-states, as per the materialist, or the eliciting of physical behaviours by subjective experiences, as per the dualist. The mind-matter gap is every bit as troubling.

It is vital to note that the inner-outer gap problem isn't simply the puzzle of how subjects outwardly appear to each other *as* structured physical objects in space-time. For this basic picture is common to all variants of panpsychism including 'anti-realist' versions that relativize physical truths to observers. What we take to be structured objects in space – Lucy's

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

brain, for instance – might exist only as the mind-dependent content of an appearance to other observers such as Jim. While this still implies an inner-outer gap of sorts, carrying its own challenges, it isn't a *mind-matter* gap that parallels those targeting materialism and dualism. The inner-outer gap problem arises for panpsychism only when there is also a mind-matter gap, with the physical world being regarded as if it had an outer, often spatial, observer-independent aspect, as demonstrated in the following passages by leading philosophers who write on panpsychism:

> what it is for physical facts p to obtain is for certain structural roles to obtain [with] ... no commitment to 'esse est percipi' ... [V]iews like this are naturally understood as versions of realism about the physical world, rather than versions of anti-realism. The physical world really exists out there, independently of our observations; it just has a surprising nature (Chalmers 2020, 354).

> The only way in which panpsychism differs from physicalism is that the basic components of the material world also involve very basic forms of consciousness, from which the more complex conscious experience of humans and other animals derives (Philip Goff 2017b).

> '[Panpsychism] leaves the universe wholly independent of our minds – except for those parts of it that are our minds. So too it leaves everything true in physics wholly in place, as remarked' (Strawson 2020, 319).

These authors clearly regard the physical world in a 'business as usual' manner, its scientific entities, laws, and discourse left external and intact. There is no mention of relativizing physical truths to the content of any appearance, human or distal. An emphatic proclamation of matter's nature as experiential, moreover, doesn't make the gap go away – it only amplifies it. Galen Strawson, for example, continues:

> [Panpsychism] is materialist by definition, and it's straight-up realist about everything that comprises what we ordinarily think of as the physical world: clouds, brains, chairs, mountains, and all the entities and qualities whose existence physics is right to recognize, quarks, say, or charge, or fields ... The ultimate, intrinsic, categorical nature of physical stuff is experience, experientiality ... experientiality is a kind of stuff: stuff = energy = experientiality. If [physical] reality is indeed spatiotemporal, then experientiality is spatiotemporal *in exactly the same way as we ordinarily suppose non-experiential stuff to be* (2020, 319–320, my italics).

We ordinarily suppose non-experiential stuff, such as clouds, brains, chairs, and mountains, to be spatiotemporal in a way that is *observer-independent*: third-personally accessible out there in the external world with real distances and locations. If this were not so there would *be* no mind-body problem. And yet experience, Strawson admits towards the end of the paper, belongs to an experiencer, in which case it is observer-dependent. How, then, has the stuff of experience suddenly become externalized in a way that puts it out there in mind-independent space, preceding the activities of physicists? That is the inner-outer gap problem.

In philosophy, the *reductio ad absurdum* strategy seeks to prove a claim by supposing its opposite to be true and then deriving from it (alongside acceptable premises) an implausible consequence that compels us to negate the starting supposition. I see the inner-outer gap problem, along with subject combination and decombination problems, as implausible

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

consequences that arise from what is, for panpsychism, an erroneous starting supposition: that there is an observer-independent physical reality with spatial locations and composition relations. Inherited from a dualist and materialist paradigm, and bolstered by how the world ordinarily presents itself to us, this supposition is the worm in the panpsychist apple that must be expunged if panpsychism is to progress beyond the (de)combination and inner-outer gap problems.

A sketch of antirealist panpsychism

Can panpsychists understand physical structures in a way that does not run into these problems? In a previous paper (Albahari 2022) I referred to the realist versions of panpsychism, that presume physical structures to be observer-independent, as 'panpsychist materialism', and the more anti-realist version to be proposed, as 'panpsychist idealism'. As Chalmers (in correspondence) thinks this muddies the terminological waters, I'll now refer to the positions as simply 'realist panpsychism' and 'anti-realist panpsychism'.

We can begin by noting what realist and anti-realist panpsychism have in common. Both versions will usually agree that the behavioural and neurological displays visible to Jim, as he watches Lucy bite into a wormy apple, are the outer appearances of inner experiences had by conscious subjects such as those we call 'Lucy' or 'Lucy's neurons'. The neurological displays will appear to Jim through patterns of cognisensory imagery that are presented to his perspective.

Where are the physical structures that characterize these objects? Here, the two positions come apart. Realist panpsychism will at least implicitly locate the structures in those observer-independent, spatio-temporal regions that we designate as 'Lucy's brain', 'neurons', 'worm', and so forth. These regions are in place prior to Jim's observation of them – their spatio-temporal structures putatively grounded in those experiences that belong to their host entities. Anti-realist panpsychism avoids the inner-outer gap problem this creates by cutting out the middle-man. There are no such spatio-temporal regions sandwiched in between the experiences of Lucy and Jim. There are only the inner experiences had by the subjects who view one another *as* spatial or mental objects. What we think of as the physical structures of the world are, moreover, grounded most immediately in patterns of cognisensory imagery that face 'towards' us, within our own human conscious fields, rather than 'away' from us, in the conscious fields of the 'distal' subjects, be they human, neural, atomic, or otherwise. The distal subjects still have a role to play in the grounding of our proximal imagery, helping explain the regularity of its observable patterns. Lucy's experience of throwing an apple to the ground shows up isomorphically to Jim *as* her throwing of an apple to the ground. But the proximal dependency of physical structures, such as those we call 'Lucy's brain', on the conscious fields of observers, such as Jim's, makes the position a recognizably anti-realist one.

How, in more detail, might anti-realist panpsychism pan out? I'll now sketch an outline that for illustrative purposes ascribes subjecthood to what we identify as 'the cosmos', along with an array of 'smaller' entities, including 'neural' or 'atomic' components.[8] I'll also suppose, as I have in my other work, that subjects are dispositional, their perspectival experiences conferring an innate power to appear to themselves and other subjects in a particular way.[9] Dispositions make the appearances seem law-like.

The cosmos is disposed to appear to humans as a vast physical plenum that contains objects in space-time obeying the laws of physics. It thus makes sense to suppose that the cosmos experiences *itself* as a sort of unified behemoth whose multiple components stand in apparent locational, ordered, and compositional relations that perhaps confer on it a sense of space, time, and ontological depth. As in our case, its component imagery – likely foreign to our own – will appear to it as objects that turn out to be the outer presentation

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

of other subjects. Items presenting to it as an organized unity, including perhaps those we call 'human' or 'atom', may be the outward appearance to it of an individual subject. Items presenting to it as mere aggregates, corresponding perhaps to such things we call 'apples' or 'mountains', might be the outward appearance to it of subject-groups whose members we know as 'cells' or 'atoms'. Of the unities, some (such as humans) may present to the cosmos as complex while others (such as those we identify as atoms) may present to it as far simpler.

Just as we will appear to the cosmos as being encased inside of itself, the cosmos appears to us as being mostly outside of our bodies. As subjects, we humans are also disposed to appear to ourselves, and each other, in various predictable ways. Within Lucy's perspective, for example, the cognisensory experience of biting into apple-with-worm (visual imagery of apple, a sense of its recognition, the intention to bite, sensations of crunch, sight of a worm) causes mental imagery of repulsion with an impulse to emit auditory imagery of 'yuck!', which in turn elicits the proprioceptive, visual, and auditory imagery of throwing an apple to the ground.

Lucy's experience registers to Jim – a similarly complex centre of subjectively based powers – as the visual imagery of her throwing an apple to the ground. Components of Lucy's reaction would be viewed by a neurologist as involving patterns of electro-chemical activity that dispose them to agree with words that include 'activity in bilateral amygdala, right inferior occipital gyrus …'. The 'neurons', supposing they are subjects, will impact on each other's imagery as they become aware of each other. They will appear collectively to Lucy as internal cognitions such as thoughts and feelings of disgust at the worm in the apple.

What appears outwardly to Jim as the arc of a falling apple will be elicited by elements of 'cosmic', and perhaps 'atomic', experience. These will show up in a physicist's conscious field as events that unfold in accordance with the Newtonian laws of physics, including units of mass, gravity, acceleration, and distance. Subatomically, the picture will appear differently again.[10]

While this is just a simple sketch of how things may hang together, in no part of the picture is there anything we would ordinarily regard as real, observer-independent objects with actual spatial locations, extensions, and distances. Lucy does not literally have a fleshy, spatially extended brain with neurons. She is not what we would ordinarily think of as a physical body in physical space that throws a physical apple to the physical ground. The information does not literally travel through space to Jim as lightwaves and soundwaves. There are only subjects appearing to one another *as if* they were (to us) spatially located, structured objects behaving in predictable, lawlike ways. The position is decidedly anti-realist, with the world and its structures wholly contained within the purview of observers. The structures we come to identify through our scientific practices will depend, most proximately and immediately, on the observational fields of human-like subjects with their particular modes of response. With no raft of observer-independent structures to explain how physical objects seem to behave, the antirealist cannot directly avail herself of the usual scientific toolbox with its reference to such concepts as space, time, gravity, mass, and momentum. Bringing these back in would usher back the inner-outer gap problem.

Is the position recognizably panpsychist? Whether subjects exist at a conventionally micro, macro, or cosmic level, the experiential life of a subject is still made from the experiences of other subjects. But not in the way construed by realist panpsychism. On a micropsychist version of that position, the experiential life of a subject (such as Lucy) is typically conceived as being made *from* the experiences that are *had* by the subjects contained inside her (such as those of her neurons). This way of envisaging matters, we saw, led straight to the combination problem with its breaching of the subject's boundaries via the sharing of experiences. On anti-realist panpsychism, the experiential life of a subject is made from *having* experiences *that target* other subjects that appear to it externally through

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

the medium of cognisensory imagery. If what we think of as various micro-entities are also subjects, then Lucy's experience is made from a presentation of other subjects (such as those we call 'atoms', 'neurons', 'worms') that collectively appear to her conscious perspective *as* apples, cognitions, and worms. No perspectival architecture is violated; the subject's privacy is preserved. *There is no combination problem*. And should what we refer to as 'the cosmos' turn out to be a subject, then just as with any other subject we will be aware of its elements via the cognisensory imagery that appears to our perspective. Through dropping the idea of the cosmos as a material entity inside of which we are spatially located, we are no longer railroaded into supposing that our conscious perspectival experience must likewise be contained inside of its conscious perspectival experience. *There is no decombination problem*.

As previously mentioned, anti-realist panpsychism does not altogether eliminate an inner-outer gap. Lucy's private experience of throwing down the apple is still disposed to show up to Jim (and other similar subjects) as roughly isomorphic patterns of behaviour and neurology that appear to him through the medium of cognisensory imagery. Explaining such isomorphism, including how subjects get to appear to one another as structured objects, is a challenge that faces *all* panpsychists, realist or anti-realist. Yet in confining all such goings-on to the medium of conscious experience, the anti-realist faces no troublesome transmutation of Lucy's subjective experience into an observer-independent raft of structural or spatio-temporal regions, such as those in her brain, that Jim then encounters and observes. The gap is not symmetrical to that facing materialism and dualism. In not bridging the disparate domains of internal mind and external matter, *there is no inner-outer gap problem*.

Some will think that renouncing the raft of observer-independent structures will sink panpsychism. I see developing new flotation devices as an exciting project. Like it or not, these are the sort of anti-realist lines along which panpsychism must develop if it is to avoid the (de)combination and inner-outer gap problems.

Objections and replies: realism fights back!

Not all realist panpsychists are happy to be charged with succumbing to the inner-outer gap problem. Galen Strawson has objected (in personal correspondence) that I have effectively strawmanned his and other panpsychist materialist positions, which include those of pioneers such as Russell, Eddington, and Strong. When Jim examines the structures of Lucy's brain or the architecture of an apple, Strawson insists, there is no troublesome transmutation of inner private experiences into a separate raft of outer observer-independent physical structures or entities. The experiences which Strawson thinks ground scientific abstractions belong to subjects that appear to us as objects such as atoms, brains, or the cosmos. These subjects have a disposition or (as he prefers the term) *power* to appear to us as elements of the familiar structured world. So while structures of these private experiences are independent of *our* observations, they are not completely observer-independent and may be isomorphic to those that we encounter in the sphere of our own experience. It is these isomorphic structures, grounded in experiences had by distal subjects, which he says anchor the subject-matter of physics, allowing statements about our external spatio-temporal world to remain literally true. In a similar vein, Chalmers has indicated (in personal correspondence) that his rendering of panpsychism does not fall prey to the inner-outer gap problem. He contends that if the experiences of subjects playing the relevant causal roles have information-structures isomorphic to those we find in the physical world, then those distal structures will be enough to secure realism about the external world. His line of reasoning is backed up by his extensive writings on the status of real-world simulations.

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

This structuralist response brings realist panpsychism into closer metaphysical alignment with anti-realist panpsychism – or at least it makes more explicit what was the case all along.[11] The world and its objects are broadly a function of how subjects appear to one another, with no separate raft of observer-independent entities to do the heavy-lifting. But realist panpsychism still departs from the anti-realist approach in two key respects. First, it claims that the physical world, with its familiar macroscopic objects and underlying structures, exists and is grounded in the experiential fields of distal subjects (such as 'cosmos' or 'atoms') as opposed to existing and being grounded most immediately within our own human experiential field. Second, it claims that these distal structures are able to serve as truthmakers for both ordinary and scientific claims about the external world. We can continue to suppose that the physical world 'really exists out there … only to have a surprising nature' (Chalmers 2020) such that scientific claims about it remain literally true. In the remainder of the article, I'll argue that this structuralist approach leaves panpsychism on unstable ground that goes against its spirit and plays back into the inner-outer gap problem. The worm still lurks at the core.

First, to the question about where a panpsychist should place what appears to us humans as the physical world, with its macroscopic objects and their micro-structures. Should it be placed most immediately in our own experiential fields, or in those of 'distal' subjects (such as 'atoms' or 'the cosmos') whose experiences are disposed to elicit object-imagery in us? We can start by using an ordinary object, the apple, as a case-study. At a macroscopic level, the fruit presents to humans as having a typical shape, taste, and crunch. It falls and rolls in a predictable way when dropped, and oxidizes when cut into. As such, the apple as we know it plays a wide range of functional and dispositional roles – disposing us to talk about apples, weigh them, cling-film them when sliced, and so on. Microscopically, the apple appears to us as having a distinctive molecular profile that disposes it to interact in a particular way with other microphysical elements.

Suppose that what externally presents to us as certain microphysical components of the apple, which I'll generically call 'atoms', are also, from inside, subjects (I'll consider a cosmic version shortly). These microsubjects with their primitive experiences, together with those of other apparently tiny players (such as 'photons'), are disposed to elicit in us the familiar apple-like imagery. But what happens when creatures such as us, with our complex set of dispositions, are taken out of the picture, leaving only the 'apples'? Assuming that an 'apple' is not itself a subject, we leave behind particular instances of monadic, primitive, private experiences and perspectives, each with their own simple structures that are elicited by taking each other in. These experiences, each of which exist only in relation to their owners, will not host the integrated structures or functional roles we associate with 'apple' any more than a pile of unlit sticks will host the role of burning down a forest. It is only when those atomic subjects are collectively registered in the conscious fields of subjects like us that the 'fire is lit,' sparking into existence the familiar imagery we associate with 'apple' and its behaviours. It is this human-like imagery that embeds the structures and roles we recognize as 'apple', so if we are going to place apples in experience, here is where to put them. For subjects that appear to us as worms, the so-called 'apple' is likely to carry a different set of roles and structures conveyed through imagery that doesn't much resemble our own. To regard them as literally apples, in this context, would be unacceptably anthropocentric.

How about the apple's structure at what we recognize to be the microscopic level? The distinctive spatially related cellular and molecular structures that we humans observe as apple-like, under microscopic conditions – including their patterns of interaction with other microphysical elements – are not concretely manifesting *as those external, spatially related patterns* in advance of our observation.[12] To insist on this would be to succumb to the inner-outer gap problem with its separate entourage of outer observer-independent structures. Prior to our observation, all we can say is that there are particular monadic instances

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

of (likely primitive) structured imagery, each manifesting to the private perspectives of what we call the various 'cells' and 'molecules'. Together, they are disposed to appear, microscopically, as 'apple-wise' to human subjects. No spatially extended apple-structures are manifesting within what is a fractured array of distal 'micro' experiences.

What if the cosmos is conscious with a global unified experience whose co-conscious structures partially reflect those of what we take to be apples? As Chalmers (2020, 365) has noted, there may be a better chance of structural isomorphism between the experiences of a cosmic subject and (what we take to be) spatio-temporal relations within the physical world than between the world and the experiences of microsubjects. However, such isomorphism around objects such as apples is likely to be highly abstract and incomplete. For without cognisensory imagery like our own, a cosmic experience that takes in what externally appears to *us* as apples is hardly likely to resemble our experience in a way that would saliently capture the particular range of roles and structures we attribute uniquely to apples, including at the microscopic level. If what we perceive as 'apple' yields two broadly isomorphic structures, why pin the apple in the mind of the cosmic subject whose experience it least resembles? Granted, (and as we'll soon see) there is a push from realist camps to get the physical structures of objects 'out there' independently of our human observation. But this brand of realism is inherited from a longstanding philosophical and scientific paradigm where we think we can get at the true nature of the world by divorcing it from all the effects of experience. If experience itself is the medium that embeds the structures and roles of objects, as it is on panpsychism, then it is far more in keeping with the spirit of the position to place those structures and roles directly in the conscious medium that conveys the objects as they characteristically present themselves to us. It is actually *more* realistic insofar as it captures a greater unity of the object's defining roles and features on both macro *and* microscopic levels.

It may be replied that this is all fine with macro-objects such as apples whose distinctive natures are entangled with our interactions. But what about the deepest physical structures of what we take to be the cosmos itself? If the cosmos is a conscious subject, then those structures that abstract away from our human sensibilities may be those we wish to isolate in understanding its ultimate physical nature. Allowing such structures to exist outside of our minds, even if they depend on the cosmic purview, may be enough to claw back the relative externality of the physical fundaments and their laws. But even accepting this, there is still good reason, as I've just argued, to suppose that at least the macroscopic world, with its familiar objects and their characteristic micro-structures, exists and is grounded most immediately within human consciousness. And since our familiar macroscopic world presents itself to us as being *outside* of our minds, its actually existing *inside* our field of consciousness would render this world a pervasive illusion, bringing the position closer to anti-realism.

The realist panpsychist may, at this juncture, recruit a line of argument put forward by David Chalmers. Incorporating also the second respect in which his structuralist approach to realist panpsychism departs from the anti-realist approach, his contention is that distal structures outside of our minds can serve as truthmakers for both ordinary and scientific claims about the world, even if nothing closely resembling the familiar macroscopic world exists outside of human experience. Proposed initially outside the context of panpsychism, the argument may be adapted to suggest that panpsychism can be realist not only about the deepest structures of physics, but also the macroscopic world itself. Its success may vindicate the existence of the external world with its structures and scientific truths, without falling into the inner-outer gap problem.

The core of Chalmers's argument can initially be found in a paper where he seeks to establish that we can be realists about the external, manifest world with its spatio-temporal objects even if we are in a Matrix-type situation. The Matrix is a movie that plays on the

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

well-known sceptical brain-in-a-vat scenario. Inputs from a computer simulation of the world are fed into the protagonist's brain, with the brain's outputs fed back into the simulation making it appear to him as if he inhabits a physical body in a physical world. The premise of the movie is that the protagonist is wrong about his situation: he is in the grip of an illusion. Rather than interacting with real objects in the real physical world, he is actually strapped into a machine undergoing a mere simulation of physical reality. While Chalmers feels the pull of this intuition, he rejects it (Chalmers 2005, 135).

The Matrix situation, he argues, relies on a combination of three metaphysical scenarios about what underlies the physical world and our minds: that physical processes are computational, our cognitive systems are separate from the physical processes with which they interact, and the physical reality was created by beings outside of physical space-time (Chalmers 2005, 136). None of these scenarios, Chalmers insists, impel us to say that chairs and tables and football teams don't exist, any more than it would if turned out that our physical world was made from quantum processes outside of space-time (which some physicists believe) or created by God. At most it would be a matter of unexpected underpinnings (Chalmers 2005, 152). What makes the simulated objects real is that their underpinnings have causal and counterfactual information structures that map perfectly onto those discoverable by physics (Chalmers 2005, 170). Elsewhere, Chalmers adds that perfect simulations of our world would satisfy a range of criteria for being real, the most central being that they are not an illusion but are 'largely as we believe them to be' (Chalmers 2022, 116).[13] Footballs are still kicked into the air, and spectators still cheer from their chairs. From such reflections Chalmers draws a wider lesson:

> One general moral is that the 'manifest image' is robust: our ordinary conception of the macroscopic world is not easily falsified by discoveries in science and metaphysics. As long as the physical world contains processes with the right sort of causal and counterfactual structure, then it will be compatible with the manifest image. Even a computer simulation has the relevant causal and counterfactual structure, as does a process in the mind of God: this is why they can support a robust external reality, despite their surprising nature (2005, 170).

This line of thought appears to transfer readily to panpsychism, something Chalmers acknowledges in *Reality +* (2022, 418). The cosmopsychist scenario would seem to most obviously approximate the 'mind of God' hypothesis, whereupon the mind of the cosmos is understood to contain conscious experiences whose structures are disposed show up to us as the physical world, whether at the micro or macroscopic level. Even if what we *take* to be the physically discoverable world is actually dependent on our minds – on a par with a simulation – its objects would be grounded not inside our minds, but in the distal structures embedded within the cosmic mind.

For Chalmers, external-world realism cuts across the conventional realist/idealist divide, insofar as it allows the causal structures of the world and its grounding to reside within the experience of a subject. Berkeley's idealism could thus qualify as realist. What matters is that the world's simulation or construction, even if manifesting in our minds, is grounded in an external (to us) process whose structures implement the exact causal and counterfactual structures of our discoverable physical world. This will include, importantly, its connection with how we experience it. In other work, Chalmers (2021) develops it as a functional analysis where physical entities, along with space, are defined by the particular functional roles that they play. The concrete nature of the reality that implements these roles and structures, he says, can be literally anything: minds, machines, cellular automata, or quantum mechanical fields.

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

For reasons just given, I would deny that the causal and counterfactual structures carried in cosmic or atomic minds *would* suffice to capture all those roles and structures that we recognize as governing the physical world, especially when characterizing familiar objects. Human input is a major determinant. But I will put this aside, and contest Chalmers's position on two other grounds.

First, I contend that the presence of an external concrete reality implementing the causal structures of our physical world is not enough to stave off anti-realist concerns that our familiar macroscopic world, at the very least, is an illusion. Chalmers is eager to press upon us the idea that a simulation, and by extension, realist panpsychist scenario, 'does not contradict any of our ordinary beliefs' even if the world as we are familiar with it is a manifestation sustained within our minds (Chalmers 2005, 152). However, it is ordinarily assumed, very deeply, that such objects as apples, chairs and clouds *in their familiar forms* exist as macroscopic entities independently of our minds – whatever their residual hidden nature might turn out to be. If the world *as it regularly presents itself to us* actually depends upon our field of consciousness – as it would do in either a computer simulation or in realist panpsychism – then there *is* a major contradiction between our ordinary beliefs about its objects and the reality behind them. There *is* a major illusion at play. Saying that the manifest world is robustly compatible with different surprising natures makes it sound, moreover, like the clash is merely between how we might theoretically expect the world to be and how it turns out rather than between what we pre-theoretically assume the world to be, and how it turns out.

My second line of response calls into question the claim that the minds of subjects, distal or otherwise, can unproblematically realize the causal and counterfactual structures of our world in a way that accords literally with our long-held scientific conceptions of those causal structures, around which their explanatory power has been classically understood. In the physical sciences it is generally assumed that *if* there are any concrete causal nodes, with powers and their effects, they will be *objects* that exist as non-mental particulars outside the minds of *all* observers – human or otherwise. Aside from effects that are had on human (or animal) minds, causal interactions are assumed to unfold directly between the objects, to be captured or explained by the law-like equations that characterize the physical sciences and which hold independently of any perspectives. The equations cover such relations as how one type of event reliably brings about another, or the patterns of composition that hold between objects at different levels of being, from the micro through to the macro and cosmic levels. This is so, even if the further nature of the causal nodes is undecided, such that their terms are treated as abstract variables, or the notion of cause and effect is open to philosophical debate. With the exception of the odd quantum-mechanical interpretation, an assumption of mind-independent, object-centred causality is baked into physical science's vast explanatory repertoire, with its concepts of space, time, particles, waves, mass, charge, momentum, inertia, velocity, and so on.

On panpsychism, the causal nodes are taken to be perspectival *subjects* with the power to bring about in themselves and other subjects the *appearances* of objects that obey natural laws. Causation can happen between different minds insofar as one subject, such as the cosmos, may exert an effect on the experiences of other subjects such as by eliciting in them the appearance of an external world. What we ordinarily assume to be non-mental objects out there in space and time, in direct causal commerce with each other, are actually mental particulars (such as sequences of imagery) which arise as the manifestations of dispositional powers belonging to different subjects, such as the cosmos or ourselves. If these manifestations – such as those we call 'billiard balls', 'atoms', or 'electrons' – are regarded in their usual way *as* causally potent non-mental particulars, in the way that science typically regards them, then contrary to how they appear they will exert no more causal power on each other than a moving shadow would have in affecting the momentum of another. In a

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

framework of subject-based causality, non-mental objects cannot thus serve as literal truth-makers for any causal roles that physics may bestow upon the objects. This is so, regardless of whether the subjective vehicles carrying the imputed structures are to be identified with the minds of human subjects or with those of distal (e.g. cosmic) observers.

The inner-outer gap problem arises when the standard scientific framework, with non-mental objects serving as the loci of causal powers, is bolted onto the panpsychist framework that has subjects serving as the seat of causal powers. The kind of causal roles that panpsychism *should* be imputing to subjects – through their power to elicit object-like appearances in themselves and one another – are displaced by the sort of causal roles that science has long imputed to non-mental objects, such as the mass role and the charge role. Subjects, under realist panpsychism, are being tasked with carrying out what are observer-independent physical events in an often spatio-temporal world – with access to the entire scientific toolbox. In being ousted from their proper causal roles, subjects are having to implement causal relations that unfold directly between non-mental particulars: of the kind that panpsychism should eschew as causally impotent!

In expounding his structuralist approach *outside* the context of panpsychism, there is ample evidence that Chalmers is availing himself of standard object-centred causality with its non-mental particulars. For example, to buttress his arguments for simulated objects being externally real, he leans heavily on analogies from the various sciences such as quantum mechanics. It is assumed that any concrete nature which embeds the relevant computational causal processes would, as with speculated quantum processes, be able to constitute the multi-levelled structures of our physical world in a way that stands clear of minds – aside from how things would appear to us (Chalmers 2022, 417).

In extending the structuralist approach to cases where causal roles are realized by the minds of subjects, Chalmers seems to be uncritically grafting object-centred causality onto the panpsychist framework. It is evident in how he and others, such as Strawson and Goff, speak of science as being business as usual, with the external, spatio-temporal world and its laws kept intact only to have a 'surprising' nature. It is evident in the way that they are curiously silent about the subject-matter of science as having to implicitly depend, as a matter of metaphysical fact, on subjects appearing to one another in various ways. It is also highly evident in the way that they model the relations between micro and macro experiences on spatial composition relations that are uncritically posited to hold between their imputed physical bearers, such as atoms, neurons and brains. It is a line of thought that has led, as I've argued, to the subject (de)combination problems.

The Russellian intrinsic nature argument indeed starts out by having us reflect that our own external physical behaviour and brain-activity – *as normally conceived* – could be implemented by our inner private mental states. This template is then applied analogously to all the outer 'behaviours' of scientific entities thought to be the causally efficacious building blocks of our external world: subatomic particles, atoms, neurons, the cosmos. After offering a promising argument that 'physical powers can only be causally efficacious in virtue of mental powers' Hedda Hassel Mørch continues that 'these mental powers relate to physical relations in a way analogous to how our will and motivation relate to our physical behaviour' (2020, 280). With no suggestion that physical behaviours are to be understood differently from usual, the hidden subjects behind the scenes are effectively shoehorned into the role of carrying out the external heavy-liftings in space, without noticing the strange mind-to-external-behaviour transmutation that would have to take place *if* the doings of physics were to be understood in the usual way.

To avoid the inner-outer gap problem, panpsychists must be anti-realists. Rather than subjects being recruited to *carry out* the externally specified roles that science demands, the so-called roles and casual structures will be *carried as* the content of subjective (re)presentations. The powers that pull the causal strings of the world will not reside in the

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

mind-independent matter of the physical sciences – their causal and compositional relations implicitly molding the grooves into which subjects must fit. The powers will reside in the experiential nature of subjects that are disposed to appear to one another *as if* there existed a mind-independent world behaving in a causally efficacious manner.

An analogy with Plato's allegory of the cave might bring the point even more sharply into focus. Residents of the cave who are confined to viewing events as shadows on the wall would be forgiven for thinking that the shadows are themselves doing the causal work. Patterns are being implemented that would suggest causality between the shadows. One shadow seems to bop another shadow on the head, causing it to run. Of course, no such causal powers really attach themselves to the shadows. They are the epiphenomenal byproducts of real goings-on behind the scenes. In a parallel fashion, under anti-realist panpsychism, it appears to us as if objects are really out there in the external world affecting one another causally *in their capacity as external objects*. Science has evolved to regard such causal relations as genuine. In reality, these happenings are like shadows on the cave-wall that merely appear to be engaging in direct causal commerce when the real causal engines are subjects hidden behind the scenes. In effect, we have the reverse of standard epiphenomenalism. Instead of the experiential states of subjects bearing the mere appearance of causal efficacy in the external world, it is the so-called components of the external world that bear the mere appearance of shouldering its causal burden.

From all these arguments, it should be evident that Strawson, Chalmers, Goff, Mørch and others *cannot* unproblematically sustain their realist conviction that panpsychism 'leaves the universe wholly independent of our minds … [with] everything true in physics wholly in place' (Strawson 2020, 319). First, panpsychism seems committed to rendering at least the familiar macroscopic world an appearance in our minds rather than something that exists, in that very form, outside of our minds. As the macroscopic world purports to exist outside of our minds, this would render it an illusion, making panpsychism, at least to that extent, anti-realist. Second, as I argued, it makes better sense for a panpsychist to ground what we take to be the world's familiar objects and their structures (such as apples) most proximally in our own human experiential fields rather than in those of distal, non-human observers (such as 'atoms' or 'the cosmos') that are disposed to appear to us as the various objects. This brings panpsychism yet closer to anti-realism. Third, whether the grounding is proximal *or* distal, subjects cannot realize the standard causal and counterfactual roles imputed to physics in a way that preserves the literal truth of physics without panpsychism falling back into the inner-outer gap problem. This again plays into the hands of anti-realism.

Conclusion: an anti-realist wolf in realist sheep's clothing?

Either panpsychists must accept that their position is committed to anti-realism or fall prey to the inner-outer gap problem. Must anti-realist panpsychists renounce the findings of physical science? Not necessarily! When it comes to doing science with its experiments, theories, and predictions, one could adopt its observer-independent paradigm as a highly useful fiction. Scientific instrumentalism could extend its scope to the point where we say: 'it is as if there were a mind-independent world with its various entities and causal laws'.[14] Yet when it comes to doing philosophy, adhering to an observer-independent paradigm can set one's modes of thinking along the wrong trajectories. This article focused on how the pervasive assumption of an observer-independent reality was the worm in the apple that led panpsychism into the traps of the (de)combination and inner-outer gap problems. But there is another way in which this worm can make trouble. In their ousting of subjects from their rightful causal roles in favour of scientific posts such as the mass role and charge role, realist panpsychists have tended to overlook the fact that they are obliged to account for *the*

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

same underlying metaphysical picture as the anti-realist panpsychist. Namely, that experiential subjects have the disposition or power to appear to one another as structured objects. What explains why all these subjects have the power to appear to one another in various lawlike ways? Are these mutual appearance dispositions just brute facts? In a point reminiscent of Kant, they cannot answer by simply appealing to the very structures that the experiences of subjects are supposed to be realizing! If this explanatory burden must be borne by *both* realist and anti-realist panpsychism, the panpsychist might do better, in the end, to go the anti-realist route, since it avoids the standard problems while remaining faithful to the spirit of the position.

Acknowledgements. Many thanks to Greg Horne, David Chalmers, and Chris Letheby for their detailed feedback on the manuscript, to David Godman for many rounds of editing suggestions, and to Galen Strawson, Philip Goff, and Andrei Buckareff for earlier exchanges.

Notes

1. While Seager (1995) coined the term 'combination problem' for micropsychism, the issue was articulated by James in 1890. I introduced the term 'decombination problem' for cosmopsychism (Albahari 2020) although the problem was again described by James (for British and American Idealism) in 1909, and other terms are also in circulation. I will focus in this article on subject versions of the (de)combination problem as they are considered the most serious.
2. For an overview of the dialectic between physicalism and dualism in relation to panpsychism see Chalmers (2016).
3. A position sometimes known as panqualityism (Chalmers 2016).
4. For more on micropsychism and cosmopsychism along with their advocates, and other variants of the (de)combination problem, see Goff et al. (2022).
5. A famous case in point is Hume's (Hume 1739–1740/1978) speaking about introspection as failing to reveal directly an observing self or subject, and yet in later thoughts admitting he cannot shake the sense that there is something 'simple' and 'continuing' behind the flux of experience. As for other subjects, we can never directly observe them in their capacity *as first-person perspectives* since they always appear to us, via our own cognisensory imagery, in the form of embodied objects from which we infer perspectivality.
6. The case of twins conjoined at the head might seem like distinct subjects that share experiences. But I think it much more likely that common, e.g., what we call 'neural', stimuli are being experienced by them similarly, but separately.
7. I present more detailed articulations of the decombination problem in Albahari (2019) and Albahari (2022).
8. This isn't cosmopsychism or micropsychism as there are no commitments to the cosmic or microsubjects being fundamental.
9. For a compelling argument for viewing all causal powers as experiential, see Mørch (2020).
10. Anti-realist panpsychism can be developed in different directions. As mentioned, the version I develop attempts to integrate it with the metaphysics of Advaita Vedanta (Albahari 2019, 2024). What we think of as objects are the outer appearances of subjects that arise as ultimately illusory entities from a substratum of pure, aperspectival consciousness.
11. In allowing that concrete particulars (e.g. minds or machines) can realize the structures, it is an epistemological rather than ontological version of structural realism (for more, see Chalmers 2022).
12. On a related point, Chalmers (2020, 361) has noted the extreme difficulty that micropsychists face in finding an experiential analogue for distance relations that physicists posit between the atomic fundaments – a worry that would extend to any distance relation on their position.
13. Five questions that Chalmers includes in the 'reality checklist' are 'Does it really exist?, Does it have causal powers?, Is it independent of our minds?, Is it as it seems?, Is it a genuine X?' (Chalmers 2022, 114).
14. Greg Horne has pointed out that the realist can also regard the standard observer-independent paradigm as a highly useful fiction, such that the mass role, for instance, 'picks out the effects of a certain experience' that exist as the content of subjective (re)presentations. But if that much is conceded, it is hard to see how the position doesn't simply collapse into anti-realist panpsychism.

References

Albahari M (2019) Perennial idealism: A mystical solution to the mind-body problem. *Philosophers' Imprint* **19**, 1–37.

https://doi.org/10.1017/S0034412525100942 Published online by Cambridge University Press

Albahari M (2020) Beyond cosmopsychism and the Great I Am: How the world might be grounded in universal (Advaitic) consciousness. In Seager W (ed.), *The Routledge Handbook of Panpsychism*. New York, NY: Routledge, 119–130.
Albahari M (2022) Panpsychism and the inner-outer gap problem. *Monist* **105**, 25–42.
Albahari M (2024) Is universal consciousness fit for ground? In Kriegel U (ed.), *Oxford Studies in Philosophy of Mind*, Vol. 4. Oxford: Oxford University Press, 3–45.
Chalmers DJ (1996) *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
Chalmers DJ (2005) The Matrix as metaphysics. In Grau C (ed.), *Philosophers Explore the Matrix*. New York, NY: Oxford University Press, 132–176.
Chalmers DJ (2016) Panpsychism and panprotopsychism. In Brüntrup G and Jaskolla L (eds), *Panpsychism: Contemporary Perspectives*. New York, NY: Oxford University Press, 19–47.
Chalmers DJ (2020) Idealism and the mind-body problem. In Seager W (ed), *The Routledge Handbook of Panpsychism*. New York, NY: Routledge, 353–373.
Chalmers DJ (2021) Finding space in a nonspatial world. In Wüthrich C, Le Bihan B and Huggett N (eds), *Philosophy beyond Spacetime: Implications from Quantum Gravity*. Oxford: Oxford University Press, 154–181.
Chalmers DJ (2022) *Reality+: Virtual Worlds and the Problems of Philosophy*. London: Allen Lane.
Goff P (2017a) *Consciousness and Fundamental Reality*. Oxford: Oxford University Press.
Goff P (2017b) The case for panpsychism. *Philosophy Now* **121**, 6–8.
Goff P (2024) How exactly does panpsychism help explain consciousness? *Journal of Consciousness Studies* **31**, 56–82.
Goff P, and Roelofs L (forthcoming) In defence of phenomenal sharing. In Bugnon M and Nida-Rümelin M, O'Conaill, D (eds), *The Phenomenology of Self-Awareness and Conscious Subjects*. New York, NY: Routledge.
Goff P, Seager W and Allen-Hermanson S (2022) Panpsychism. In *The Stanford Encyclopedia of Philosophy*. (Summer 2022 Edition) Zalta EN (ed.). https://plato.stanford.edu/archives/sum2022/entries/panpsychism/.
Hawking S (1988) *A Brief History of Time: From the Big Bang to Black Holes*. New York, NY: Bantam Books.
Hume D (1739–1740/1978) *A Treatise of Human Nature*. Selby-Bigge LA, and Nidditch PH (eds). Oxford: Clarendon Press it should say Edited by L. A. Selby-Bigge; 2nd ed., revised with variant readings by P. H. Nidditch.
James W (1890) *The Principles of Psychology*. London: Macmillan & Co.
James W (1909/1977) *A Pluralistic Universe*. Cambridge, MA: Harvard University Press.
Mørch HH (2020) The argument for panpsychism from experience of causation. In Seager W (ed.), *The Routledge Handbook of Panpsychism*. New York, NY: Routledge, 269–284.
Russell B (1927) *The Analysis of Matter*. London: Kegan Paul.
Schaffer J (2010) Monism: The priority of the whole. *Philosophical Review* **119**, 31–76.
Seager WE (1995) Consciousness, information and panpsychism. *Journal of Consciousness Studies* **2**, 272–288.
Strawson G (2020) What does 'physical' mean? A prolegomenon to physicalist panpsychism. In Seager W (ed.), *The Routledge Handbook of Panpsychism*. New York, NY: Routledge, 317–339.

Cite this article: Albahari M (2025) Does panpsychism entail anti-realism? The worm in the panpsychist apple. *Religious Studies*, 1–19. https://doi.org/10.1017/S0034412525100942