

Bernardo Kastrup

The Universe in Consciousness

***Abstract:** I propose an idealist ontology that makes sense of reality in a more parsimonious and empirically rigorous manner than mainstream physicalism, bottom-up panpsychism, and cosmopsychism. The proposed ontology also offers more explanatory power than these three alternatives, in that it does not fall prey to the hard problem of consciousness, the combination problem, or the decombination problem, respectively. It can be summarized as follows: there is only cosmic consciousness. We, as well as all other living organisms, are but dissociated alters of cosmic consciousness, surrounded by its thoughts. The inanimate world we see around us is the extrinsic appearance of these thoughts. The living organisms we share the world with are the extrinsic appearances of other dissociated alters.*

1. Brief Introduction

This paper seeks to articulate an ontology that overcomes the principal limitations of the most popular alternatives. The first half of the paper comprises a detailed analysis of relevant literature, highlighting what advances have been made and what problems have been created or left unsolved by recent developments in analytic philosophy. In the second half, starting from what I consider to be the most promising current platform, I propose an idealist framework that may open viable new avenues for addressing the key questions left unanswered by this current platform. At the end, I hope to offer a coherent view of the nature of reality that accounts for all relevant facts without incurring any fundamental problem.

Correspondence:
Email: bernardo@bernardokastrup.com

2. The Mainstream Physicalist Ontology and its Problems

The mainstream ontology of physicalism posits that reality is constituted by irreducible entities — which, like Galen Strawson (2006, p. 9), I shall call ‘ultimates’ — outside and independent of phenomenal consciousness. These ultimates, in and of themselves, do not instantiate phenomenal properties: there is nothing it is like to be an ultimate, the capacity for experience emerging only at the level of complex arrangements of ultimates. They are also sometimes held to lack objective qualities: in and of themselves, ultimates may have no colour, flavour, smell, etc. Indeed, according to mainstream physicalism qualities may exist only in the phenomenal field of the experimenter, which in turn is a product of the operation of a sufficiently complex nervous system. It is the specific arrangement of ultimates in a nervous system that, allegedly, somehow constitutes or generates its phenomenal properties.

The key problem of mainstream physicalism centres on how our subjective experience of qualities — what it is like to feel the warmth of fire, the redness of an apple, the bitterness of disappointment, etc. — can arise from mere arrangements of ultimates. These ultimates do possess abstract relational properties such as mass, spin, momentum, and charge, but there is nothing about mass, spin, momentum, or charge, or the relative positions and interactions across ultimates, in terms of which one could deduce what the warmth of fire, the redness of an apple, or the bitterness of disappointment feel like, subjectively. As long as they fit with the broadly observed correlations between neural activity and reported experience, mappings between these two domains are entirely arbitrary: in principle, it is as (in)valid to state that spin up constitutes or generates the phenomenal property ‘coldness’ and spin down ‘warmth’ as it is to say the exact opposite. There is nothing intrinsic about spin — or about any other property of ultimates or arrangements thereof — that would allow us to make the distinction.

This central — and arguably insoluble — problem has been referred to by different names, such as the ‘explanatory gap’ (Levine, 1983) and, more recently, the ‘hard problem of consciousness’ (Chalmers, 1996; 2003): the qualities of experience are irreducible to the observable parameters of physical arrangements — whatever the arrangement may be — in the sense that it is impossible even in principle to deduce those qualities from these parameters. More

generally, the argument here is that there is no *entailment* from facts about ultimates to facts about experience: there is no fact about ultimates that implies *a priori* a fact about experience.

Greg Rosenberg (2004, pp. 13–30) articulated what is perhaps the best refutation of entailment from facts about ultimates to facts about experience. His argument begins with the recognition that all facts about ultimates are merely patterns of bare differences. This echoes Bertrand Russell's point (2007) that science can only characterize things and phenomena in terms of how they differ from other things and phenomena. For instance, an ultimate with positive electric charge is characterized in terms of how its relevant behaviour differs from that of a negatively charged ultimate. Charge is thus a relational property defined on the basis of bare differences. Nothing can be scientifically stated about what a charge, in and of itself, intrinsically is. The same can be argued about all other facts about ultimates.

Rosenberg then proceeds to show that facts about experience — phenomenal properties — cannot be entailed by patterns of bare differences, even though qualitative differences between experiences can admittedly instantiate a structure of bare differences. Therefore, phenomenal properties cannot be reduced to facts about ultimates. Allow me to unpack this.

There are qualitative differences across our experiences of various colours: what it is like to see yellow is different from what it is like to see red. These qualitative differences can even be graded along relevant dimensions: the qualitative difference between seeing yellow and red seems bigger than the qualitative difference between seeing yellow and orange. If one were to assign a number to represent each of these degrees of difference, one could abstract out a purely quantitative — that is, bare — difference structure from the experiences of seeing various colours. However, that a bare difference structure can be abstracted out from phenomenal properties does not imply that phenomenal properties are entailed by bare difference structures. Maintaining so inverts the logic of the situation: it is phenomenal properties that ground bare difference structures in the first place.

To bring this point home, Rosenberg offers the following thought experiment: imagine a field of tightly packed yellow and red dots. If one observes this field from a sufficient distance, one sees the colour orange. It could then be argued that the phenomenal property 'orange' arises from a pattern of bare differences associated with the delta in wavelength between yellow and red photons, as well as the relative size and distribution of the dots. However, if one were to choose

another pair of colours with the same delta in wavelength — say, yellow and green — and otherwise maintain the same relative structure of dots, a phenomenal property different from ‘orange’ would result. In other words, the same pattern of bare differences would yield a different phenomenal property. Hence, phenomenal properties are not entailed by patterns of bare differences and cannot be reduced to properties and arrangements of ultimates.

This and other arguments along similar lines render mainstream physicalism arguably untenable.

3. Consciousness as an Irreducible Property of Matter

At least since the time of René Descartes, the most recognizable alternative to physicalism has been ‘substance dualism’: if one cannot reduce phenomenal properties to physical elements, then the phenomenal and the physical may be two distinct, fundamental ontological classes. There are different versions of substance dualism, but the most intuitive one is arguably ‘interactionism’: since phenomenal events seem to cause physical events (as in when felt pain causes me to move my arm) and vice versa (as in when a needle piercing my arm causes me to feel pain), then the phenomenal and the physical must be causally connected. However, a problem with interactionism is summarized by Chalmers (2016b, p. 23): if the physical domain is causally closed — as it seems to be in so far as we have been able to ascertain through the scientific method — then causal influences we intuitively attribute to the phenomenal domain must ultimately be, in fact, physical. There is arguably no place for phenomenal properties in the causal nexus. Possible dualist answers to this have been proposed but, as acknowledged by Chalmers himself (who admits to sympathy towards dualism), ‘there is at least a *prima facie* case against dualism here’ (*ibid.*, p. 24). Chalmers then posits an ‘Hegelian synthesis’ (*ibid.*) between mainstream physicalism and substance dualism, in the form of the notion that ultimates themselves may be *fundamentally* conscious.

Indeed, under mainstream physicalism, ultimates are elementary subatomic particles — quarks, leptons, gauge bosons, and scalar boson(s) — with certain fundamental properties. These properties are relational and abstract, such as mass, charge, spin, and momentum. Mainstream physicalism’s key problem, as we have seen, is its inability to account for phenomenal properties. So the most

straightforward way out is to posit that at least some elementary particles *also* have fundamental phenomenal properties. In Strawson's words, 'Assuming, then, that there is a plurality of physical ultimates, some of them at least must be intrinsically experiential, intrinsically experience-involving' (2006, p. 24).

I shall call these experiencing elementary particles 'phenomenal ultimates'. I shall also generally refer to the broad ontological outlook described above as 'bottom-up panpsychism', even though I am aware that there are many variations of it that would be better discriminated from one another (such as 'panexperientialism', 'constitutive micro-psychism', 'panprotopsychism', 'deferential monadic panpsychism', etc.). Be that as it may, the key general idea here is that, by positing phenomenal properties to be fundamental, bottom-up panpsychism evades the need to reduce these properties and thereby avoids the hard problem altogether. Moreover, bottom-up panpsychism places these new fundamental properties seamlessly alongside existing abstract relational properties, as the categorical basis of the latter. This neatly integrates phenomenal properties in the framework of scientific thinking, for they now occupy a proper place within the causal nexus.

To see why this seemingly elegant approach nonetheless fails, notice that, according to bottom-up panpsychism, the unitary phenomenal life of a human being is supposedly constituted by micro-level phenomenal parts. At some point in the remote past phenomenal ultimates

organized into increasingly complex forms, both experiential and non-experiential, by many processes including evolution by natural selection. And just as there was spectacular enlargement and fine-tuning of non-experiential forms (the bodies of living things), so too there was spectacular enlargement and fine-tuning of experiential forms. (*ibid.*, p. 27)

However, the idea that micro-level phenomenal states can combine to form unitary macro-level phenomenal states is arguably incoherent. It leads to a variety of 'combination problems' (Chalmers, 2016a), at least one of which is arguably as insoluble as the hard problem itself (Carruthers and Schechter, 2006; Goff, 2006; 2009).

The best argument against bottom-up panpsychism is perhaps Sam Coleman's (2014). As bottom-up panpsychists themselves seem to agree, "experience is impossible without an experiencer," a *subject of experience*' (Strawson, 2006, p. 26, emphasis added). Therefore, bottom-up panpsychism implies that each phenomenal ultimate, by virtue of bearing phenomenal properties, instantiates a micro-level

subject. Moreover, it implies that macro-level subjects with a seemingly unitary perspective, such as you and me, must somehow arise through some form of bottom-up combination of micro-level subjects. This is called the ‘subject combination problem’.

Coleman connects subjectivity with the presence of a perspective, or *point of view*:

That a given subject has a particular phenomenological point of view can be taken as saying that there exists a discrete ‘sphere’ of conscious experiential goings-on corresponding to this subject, with regard to which other subjects are distinct in respect of the phenomenal qualities they experience, *and they have no direct (i.e. experiential) access to the qualitative field enjoyed by the first subject*. A subject, then, can be thought of as a *point of view* annexed to a *private* qualitative field. (Coleman, 2014, p. 30, emphasis added)

Notice Coleman’s emphasis on the *private* character of the qualitative field annexed to a subject. I shall return to this point later.

Bottom-up panpsychism attempts to model the combination of phenomenal states after the way ultimates combine physico-chemically. After all, the force and appeal of its argument rests on the analogous treatment of phenomenal properties and standard physical properties such as mass, spin, and charge. Therefore, Coleman also makes explicit what *combination* means in this physico-chemical sense:

Combination, thus, is the formation of a whole from components where the components continue to exist in the whole, but are intrinsically altered by combining with one another. (*ibid.*)

For instance, an oxygen and two hydrogen atoms *combine* to form a water molecule: they become intrinsically altered in the process of forming covalent bonds with one another, but continue nonetheless to exist in the resulting molecule.

In this framework, bottom-up panpsychism implies that the private point of view of each phenomenal ultimate that constitutes you becomes intrinsically altered in the process of combining to form the private point of view you enjoy right now — that is, your ‘unique experiential portal to reality’, as put by Itay Shani (2015, p. 399). *But each must nonetheless continue to exist in you*, just as quarks continue to exist in protons, protons continue to exist in oxygen atoms, and oxygen atoms continue to exist in water molecules.

However, Coleman argues, ‘points of view cannot combine’ in this manner (2014, p. 32). If a first constituent lower-level subject sees,

say, only blue, and a second sees only red, then only the *qualitative contents* of their respective perspectives can conceivably survive — possibly in altered form — as combined ingredients of the resulting higher-level subject's phenomenal field (e.g. if the latter sees purple). But the original constituent *points of view* cannot survive, for they entail seeing *only* red and *only* blue, respectively. Since the resulting higher-level subject has, *ex hypothesi*, a single compound portal to reality, it cannot *both* see only red *and* only blue. At least one of the constituent lower-level points of view will thus necessarily disappear — in fact, *both* will disappear if the higher-level subject sees purple — which is not consistent with combination in the physico-chemical sense.

One may argue that what happens instead is that the phenomenal state of the higher-level subject 'is a novel state which in some way "absorbs" or supersedes the mental states of the constituents' (Seager, 2010, p. 179). In this so-called 'combinatorial infusion' (*ibid.*) scenario, the lower-level points of view cease to exist in the process of forming the compound higher-level one. By parting with combination in the physico-chemical sense, this scenario negates much of the force and appeal of the bottom-up panpsychist argument. But the panpsychist can be spared this regret, for — as Coleman argues — the scenario does not work anyway.

Coleman's reasoning is that, to avoid the appeal to magic entailed by brute or strong emergence, 'lower-level properties must contribute to their novel product *in virtue of their metaphysical nature*, or, otherwise put, *while remaining true to what they are*' (2014, p. 35, original emphasis). But 'a set of points of view have nothing to contribute *as such* to a single, unified successor point of view. Their essential property defines them against it: in so far as they are points of view they are experientially distinct and isolated' (*ibid.*, p. 37, original emphasis). So the resulting higher-level point of view cannot be explicated in terms of the lower-level constituent points of view.

In conclusion, bottom-up panpsychism fails because there is no explicit and coherent way to ground the existence of macro-level subjects in micro-level phenomenal ultimates. Subject combination arguably requires — just as mainstream physicalism does — the appeal to magic entailed by brute or strong emergence. Yet, it was precisely this requirement that, in the case of mainstream physicalism, motivated the conception of bottom-up panpsychism as an alternative in the first place.

4. The Questionable Logical Bridge in Bottom-Up Panpsychism

Bottom-up panpsychism is motivated by the idea that, since physics only models the *behaviour* of physical entities and says nothing about their *intrinsic nature* (Russell, 2007), phenomenal consciousness may be this intrinsic nature. This is eminently reasonable, since the only physical entity we are acquainted with ‘from within’ is our own nervous system, whose intrinsic nature surely seems to be phenomenal (Eddington, 1928). But bottom-up panpsychism then makes an *extra* claim: that phenomenal consciousness has the same *fragmented structure* that matter has on the screen of perception. In other words, since our body is constituted by myriad elementary particles in so far as we can perceive it, our phenomenal inner-life must itself be constituted by micro-level phenomenal parts — or so the argument goes.

This extra claim rests on a questionable logical bridge: it attributes to *that which experiences* a structure discernible only in the *experience itself*. Allow me to elaborate.

The concept of elementary particles — ultimates — arises from experiments whose outcomes are accessible to us only in the form of conscious perception (even when delicate instrumentation is used, the output of this instrumentation is only available to us as conscious perception). Such experiments show that the images on the screen of perception can be divided up into ever-smaller elements, until we reach a limit. At this limit, we find the smallest discernible constituents of the images, which are thus akin to pixels. As such, ultimates are the ‘pixels’ of *experience*, not necessarily of the *experiencer*. The latter simply does not follow from the former.

Therefore, that human bodies are made of elementary particles does not necessarily say anything about the structure of the *experiencer*: a human body is itself an image on the screen of perception, and so will necessarily be ‘pixelated’ in so far as it is perceived. Such pixelation reflects the idiosyncrasies of *the screen of perception*, not necessarily the structure of the human *subject* itself. As an analogy, the pixelated image of a person on a television screen reflects the idiosyncrasies of the television screen; it does not mean that the person herself is made up of pixels.

As suggestive as it may be, the hypothesis that phenomenal consciousness is the intrinsic nature of the physical does not imply that the fragmented structure of matter on the screen of perception is the fundamental structure of phenomenal consciousness itself.

5. What Counts as a Fundamental Concrete Entity?

We have seen in the previous section that elementary particles are the building blocks or ‘pixels’ of what is perceived, not necessarily of the subject that perceives. But we can ask a yet deeper question: are elementary particles fundamental concrete entities on their own merit? Both mainstream physicalism and bottom-up panpsychism, in taking ultimates to be the discrete building blocks of nature, seem to assume so.

There are, however, strong reasons to believe that at least the entire inanimate universe is one integrated whole without ultimate parts. Jonathan Schaffer, for instance, points out that

physically, there is good evidence that the cosmos forms an entangled system and good reason to treat entangled systems as irreducible wholes. Modally, mereology allows for the possibility of *atomless gunk*, with no ultimate parts for the pluralist to invoke as the ground of being. (2010, p. 32, original emphasis)

Terry Horgan and Matjaž Potrč (2000) also contended that only the universe as a whole can be considered a concrete entity on its own merit, which they called the ‘bobject’.

The physical substantiation for this line of thought is not recent. As early as in the 1930s, John von Neumann (1996) reasoned that, when two inanimate quantum systems interact, no measurement is actually performed but, instead, the two systems become entangled with one another, forming an indivisible whole. If the resulting whole then interacts with a third system, they, too, become entangled, forming a new and larger whole; and so forth. These are the so-called ‘von Neumann chains’ and, since everything in the universe ultimately is a quantum system, the entire inanimate universe must constitute one single von Neumann chain — that is, one indivisible whole (von Neumann also remarked that observation by a conscious, living human being clearly breaks the chain, since living humans demonstrably *can* perform a quantum measurement. Therefore, conscious living beings must be left out of the present argument).

The implication is that, physically, there are arguably no such things as fundamental microscopic ultimates. Although this may violate popular assumptions and intuitions, it also points the way to a third avenue of enquiry that holds some promise as an alternative to both mainstream physicalism and bottom-up panpsychism.

6. The Whole Universe as a Unitary Conscious Entity

The idea that the (inanimate) universe may be an indivisible whole has proven tempting to those seeking an alternative to bottom-up panpsychism, so to avoid the subject combination problem: they posit that ‘*the cosmos as a whole is the only ontological ultimate there is, and that it is conscious*’ (Shani, 2015, p. 408, original emphasis). With this, there is no longer any need to explicate how lower-level subjects combine to form higher-level subjects, for the highest-possible-level subject is already the starting point.

This general outlook is called ‘cosmopsychism’ (Mathews, 2011; Jaskolla and Buck, 2012; Shani, 2015; Nagasawa and Wager, 2016). The seminal insight that freed cosmopsychism from the limitations of bottom-up panpsychism was arguably that of Freya Mathews (2011): she realized that, even under the hypothesis that phenomenal consciousness is the intrinsic aspect of the physical, there is no need to attribute the fragmented structure of matter to phenomenal consciousness itself. In her words, ‘an extension of subjectivity to physical reality generally [i.e. “force fields and even space itself”], rather than its restriction merely to matter, does seem to be required’ (*ibid.*, p. 144).

Now the problem cosmopsychists face is the ‘decombination problem’ (also called the ‘decomposition problem’ in Chalmers, 2016a): how do seemingly separate *lower*-level subjects — which, from now on, I shall follow Shani (2015, p. 415) in referring to as ‘relative subjects’ — form within the conscious cosmos? To paraphrase Coleman (2014, p. 30), how do they acquire their *private* point of view, whose associated qualitative field other relative subjects have no direct — that is, experiential — access to? After all, I cannot read your thoughts and, presumably, neither can you mine.

Before we address this problem, however, notice that there are at least two possible interpretations of cosmopsychism. The first one sticks to the bottom-up panpsychist view that a phenomenal ultimate has both phenomenal *and non-phenomenal* properties. This way, whereas it takes the cosmos as a whole to be the sole phenomenal ultimate there is, this interpretation grants that the abstract relational properties of the cosmos are not phenomenal. For this reason, I shall call this interpretation ‘dual-aspect cosmopsychism’. According to it, the intrinsic aspect of the cosmos is phenomenal, but its extrinsic aspect — the physical structure we can objectively measure in a scientific sense — is non-phenomenal and circumscribes the cosmos’s

phenomenal field. In a sense, the extrinsic, physical aspect of the cosmos *bears* phenomenality within in.

Another interpretation of cosmopsychism entails that *the sole ontological primitive there is is cosmic phenomenal consciousness* — or simply ‘cosmic consciousness’ for ease of reference. Nothing exists outside or independent of cosmic consciousness. As such, under this interpretation one should say that the cosmos is *constituted by* phenomenality, as opposed to *bearing* phenomenality. In other words, here the perceivable cosmos is *in consciousness*, as opposed to being *conscious*.

The latter interpretation is Shani’s (2015) position. Indeed, according to him the external, physical aspect of the cosmos is ‘its *appearance* as an exterior complement to... subjective realities’ (*ibid.*, p. 412, emphasis added). Appearances are, of course, phenomenal in nature. I shall thus call this interpretation ‘idealist cosmopsychism’, since its reduction base is purely phenomenal.

Shani does still postulate a duality in cosmic consciousness to account for the clear qualitative differences between the outer world we, as relative subjects, perceive and measure and the inner world of our thoughts and feelings. He calls it the ‘lateral duality principle’ (*ibid.*, p. 410) and describes it thus:

[Cosmic consciousness] exemplifies a dual nature: it has a *concealed* (or *enfolded*, or *implicit*) side to its being, as well as a *revealed* (or *unfolded*, or *explicit*) side; the former is an intrinsic dynamic domain of creative activity, while the latter is identified as the outer, observable expression of that activity. (*ibid.*, original emphasis)

What is important to emphasize, though, is that this duality does not entail or imply two distinct ontological classes. Everything is still phenomenal.

Now, one must ultimately ground the revealed side of the cosmos in its concealed side, not only to eliminate what would otherwise be an arbitrary boundary, but also to accommodate the empirically undeniable causal links between the revealed order of the physical world we perceive and the concealed order of thoughts and feelings. After all, revealed physical things and phenomena — think of psychoactive drugs, bodily trauma, electromagnetic fields, etc. — causally affect our concealed thoughts and feelings. Causal links operating the other way around are also undeniable: our thoughts and feelings can lead to physical manifestations in the form of bodily behaviours. If the revealed order were not grounded in the concealed, but constituted a separate phenomenal domain instead, how could these cross-

influences take place? Indeed, Shani acknowledges as much when he writes that ‘the revealed order of reality is grounded in the concealed’ (*ibid.*, p. 416).

Yet, Shani is not explicit in regards to *how* this grounding works. He states that the physical world we perceive is the way the structural patterns of the creative activity of cosmic consciousness are represented in relative subjects, such as you and me (*ibid.*, p. 412). This is fair enough as far as it goes, but what is the mechanism of representation whereby concealed phenomenal activity translates into revealed order from the perspective of relative subjects? How does the formation of a relative subject lead to such a significant qualitative transition as the representation of thoughts and feelings (the concealed order) in the form of perception (the revealed order)?

To tackle the decombination problem, Shani posits that the conscious perspective or point of view of each relative subject has both a specific and a generic character (*ibid.*, p. 423). Since a relative subject corresponds to a segment of cosmic consciousness, its specific character is derived from the *local pattern of phenomenal activity* taking place in that segment. Its generic character, in turn, is derived from the *intrinsically subjective, perspectival nature* of cosmic consciousness as a whole. Let me unpack this.

Shani posits two intrinsic features of cosmic consciousness as constituents of the generic character of each relative subject: *sentience* and *core-subjectivity* (*ibid.*, p. 426). In other words, each relative subject is phenomenally conscious by virtue of the fact that cosmic consciousness is itself intrinsically capable of experience. Also, each relative subject has ‘ipseity, or I-ness, by which is meant an implicit sense of self which serves as the dative... of experience, namely, as *that to whom things are given*, or disclosed, from a perspective’ (*ibid.*, original emphasis). The claim is then that the sense of I-ness of each relative subject *is* the sense of I-ness intrinsic to cosmic consciousness as a whole. One could argue that sentience and core-subjectivity, so defined, are inextricable from one another. But even in this case, it is still useful to distinguish between these two cognitively salient aspects of what would admittedly be a single intrinsic feature of cosmic consciousness. So I shall continue to speak of sentience and core-subjectivity.

In summary, according to Shani a relative subject is grounded, on the one hand, in the intrinsic sentience and core-subjectivity of cosmic consciousness as a whole and, on the other hand, in the local patterns of phenomenal activity taking place in the particular segment of

cosmic consciousness associated with the relative subject. The question now is: what are these local patterns of phenomenal activity that give rise to a *private* qualitative field, *inaccessible to other relative subjects*, as required by Coleman (2014)?

Shani posits that the smallest *cohesive* elements of nature correspond to the revealed appearance of micro-level relative subjects (2015, pp. 415–16). In other words, *he returns to the bottom-up panpsychist view that elementary particles are subjects*. Shani motivates this with a metaphor:

A relative [subject] is a ‘vortex’ surging from the oceanic background [of cosmic consciousness]. It is a cohesive system. (*ibid.*)

He uses the image of a vortex to refer to localization of phenomenal activity.

Consider... the *most elementary* ‘vortices’. [Their corresponding] localization process consists, then, in the intensification of experience, as well as in the concentration of focus, within limited and relatively well-defined boundaries... which serves to separate the system’s inner reality from the inner reality of the ocean surrounding it... The result is an individual self (*however primitive*) engulfed in its own experiences and concerns while being ignorant of the deeper layers which bind it to the ground of all things... [*The theory implies that simple* [vortices] *are veritable subjects*. (*ibid.*, p. 418, emphasis added)

Having effectively returned to the idea of micro-level phenomenal ultimates, Shani then argues that macro-level relative subjects, such as you and me, *are formed by micro-level relative subjects coming together*. The rather technical core of his argument — which I shall not reproduce here, for it is not relevant to this paper — is that, by grounding the micro-level relative subjects in cosmic consciousness, he circumvents Coleman’s (2014) attack on bottom-up panpsychism.

Even if the latter point is valid — and I have no reason to believe otherwise — I see multiple problems with this move. For one, once one starts from *cosmic* consciousness, it seems unnecessary and rather convoluted to descend all the way down to micro-level subjects, just to turn around again and go up to macro-level subjects. The only motivation I see for doing so is the arguably flawed notion, discussed earlier, that the ‘pixels’ discernible on the screen of perception must be the building blocks of the *experiencer*, as opposed to the *experience*. By making a concession to this physicalist intuition, Shani forces *two* problems upon himself: he has to explain (a) how the cosmic subject seemingly breaks up into myriad micro-level relative

subjects, and then (b) how these micro-level relative subjects come together again to form macro-level relative subjects.

Moreover, recall that, as per Coleman's definition, subjects entail 'a point of view annexed to a private qualitative field' (Coleman, 2014, p. 30). So micro-level relative subjects must have *private* phenomenal fields *inaccessible to other subjects*. To tackle the decombination problem, one must explain how these private fields form within the ocean of cosmic consciousness. But Shani seems to address this only in a vague, tangential manner. For starters, it is unclear how or why a mere 'localization process' in the ocean of cosmic consciousness would lead to local 'intensification of experience' and 'concentration of focus' (Shani, 2015, p. 418). But even if we grant that it somehow does, a 'concentration of focus within limited and relatively well-defined boundaries' does not seem sufficient 'to separate the system's inner reality from the inner reality of the ocean surrounding it' (*ibid.*). By way of analogy, while my visual focus right now rests on the characters I am writing, I am not unaware of, or separate from, the contents of my peripheral vision; I still have direct — that is, experiential — access to them. Moreover, it is hard to imagine that an electron could have sufficiently rich phenomenal properties to become 'engulfed in its own experiences and concerns' (*ibid.*). It seems highly unlikely that there is enough cognitive complexity — if there can be cognition at all — at that microscopic level to justify such an appeal to mere self-absorption as the mechanism behind the separation of the electron's inner reality.

I do think Shani is on the correct general path here, but a more elaborate, explicit, and precise case, with stronger *empirical* substantiation, seems to be necessary to tackle the decombination problem.

7. The Key Questions to Be Answered

The principle of parsimony implies that, of the two interpretations of cosmopsychism discussed above, idealist cosmopsychism is more economical and, therefore, should be preferred *if it can account for all relevant facts*. I shall thus take idealist cosmopsychism as my starting point and then attempt to address each of its problems and limitations. The goal is to account for all relevant facts with cosmic consciousness alone in the reduction base. Because I do not feel the need to invent new names for ideas that have historically established names, I shall call the resulting ontology simply *idealism*.

Specifically, here are the key problems of, and questions not sufficiently or explicitly addressed by, idealist cosmopsychism that I now set out to tackle:

- a) *Grounding experience in cosmic consciousness*: how do myriad, ephemeral experiential qualities arise in one enduring cosmic consciousness?
- b) *The decombination problem*: how do private phenomenal fields form within cosmic consciousness? Why can I not read your thoughts by simply shifting the focus of my attention?
- c) *Reducing perception*: how can the revealed order of nature (the physical world we measure) be explained in terms of its concealed order (its underlying thoughts)? Why are the respective qualities so different?
- d) *Explaining the correlations between brain function and inner experience*: if brain function does not constitute or generate phenomenality, why do they correlate so well?
- e) *Explaining a seemingly shared, autonomous world*: if the world is imagined in consciousness, how can we all be imagining essentially the same world outside the control of our personal volition?

8. Experiences as Excitations of Cosmic Consciousness

The first step is to clarify the relationship between cosmic consciousness and experience. After all, the two are not interchangeable: cosmic consciousness is, *ex hypothesi*, something relatively enduring and stable, whereas experiences are relatively ephemeral and dynamic. Yet, idealism posits that cosmic consciousness is nature's *sole* ontological primitive, so how does the variety and dynamism of experience come into the picture?

I submit that (a) experiences are *patterns of self-excitation* of cosmic consciousness and that (b) cosmic consciousness has the inherent disposition to self-excitation. As such, experiences are not ontologically distinct from cosmic consciousness, just as a dance is not distinct from the dancer. There is nothing to a dance but the dancer in motion. In an analogous way, there is nothing to experience but cosmic consciousness 'in motion'.

Particular experiences correspond to *particular patterns* of self-excitation of cosmic consciousness, just as particular choreographies correspond to particular patterns of self-excitation of the dancer. These patterns can evolve in time and differ across different segments

of cosmic consciousness. It is the variety and dynamics of excitations across the underlying ‘medium’ that lead to different experiential qualities. (One must be careful at this point: by referring to cosmic consciousness as a ‘medium’ I may appear to be objectifying it. Language forces me into this dilemma. But cosmic consciousness is subjectivity itself, not an object.) This way, even if the ‘medium’ is eternal and immutable, its self-excitations can come and go in myriad patterns.

This notion is entirely analogous to, and consistent with, how modern physics attempts to reduce the variety and dynamics of natural phenomena to an enduring primary substrate: quantum field theory, for instance, posits that all fundamental particles are particular modes of self-excitation of a quantum field, which is inherently disposed to self-excitation. Superstring theories posit essentially the same, but now the self-excited substrate is hyper-dimensional strings. Finally, according to M-theory the patterns of nature consist of modes of self-excitation of a hyper-dimensional membrane. Idealism, as I am formulating it here, essentially entails porting the evolving mathematical apparatus of modern physics to cosmic consciousness itself, as opposed to an abstract conceptual object. This should require but a straightforward and seamless transposition, implying no loss of predictive power.

9. Tackling the Decombination Problem

Cosmic consciousness comprises a variety of *phenomenal contents* — experiences, patterns of self-excitation — such as thoughts and feelings. If we take the human psyche as a representative sample of how cosmic consciousness operates — which is the best we can do, really — we can infer that, ordinarily, these phenomenal contents are *internally integrated* through *cognitive associations*: a feeling evokes an abstract idea, which triggers a memory, which inspires a thought, etc. These associations are *logical*, in the sense that, for instance, the memory inspires the thought because of a certain implicit logic linking the two. Ordinary phenomenal activity in cosmic consciousness can thus be modelled as a connected directed graph. See Figure 1a. Each vertex in the graph represents a particular phenomenal content and each edge a cognitive association logically linking contents together. Every phenomenal content in the graph of Figure 1a can be reached from any other phenomenal content through a chain of cognitive associations.




Figure 1. A connected graph (a) illustrating normal integration of phenomenal contents, and a disconnected graph (b) illustrating dissociation and the corresponding formation of an alter (inner subgraph in grey).

Each vertex in Figure 1 represents a particular pattern of self-excitation of cosmic consciousness. Each edge represents thus an association between two patterns of self-excitation, each pattern with its particular constituent harmonics. When the two patterns of self-excitation are concurrently present — that is, when the two associated phenomenal contents are experienced together — the association can be seen as a combination of the respective harmonics, like in a musical chord wherein multiple notes are played at the same time. When the association unfolds in temporal sequence — as e.g. in the case of a thought that fades away to make room for the experience of the memory it evokes — it can be visualized as a transition from the first to the second pattern of self-excitation, like notes played in sequence in a melody.

However, we know from the psychiatric literature that sometimes ‘a disruption of and/or discontinuity in the normal integration’ of phenomenal contents can occur in the human psyche (Black and Grant, 2014, p. 191). This is called *dissociation* and is well recognized in psychiatry today (American Psychiatric Association, 2013). Dissociation entails that some phenomenal contents cease to be able to evoke others. A person suffering from a particularly severe form of dissociation, called Dissociative Identity Disorder (DID), exhibits multiple ‘discrete centers of self-awareness’ (Braude, 1995, p. 67) called *alters*. Each alter corresponds thus to a particular segment of the psychic space wherein it forms.

Dissociation can be visualized as what happens when the graph in Figure 1a becomes disconnected, such as shown in Figure 1b. Some phenomenal contents can then no longer be reached from others. The

inner subgraph is thus a representation of an alter, corresponding to a particular segment of the originally integrated psychic space.

There is compelling empirical evidence that different alters can remain concurrently conscious. In Morton Prince's well-known study of the 'Miss Beauchamp' case of DID, one of the alters 'was a co-conscious personality in a deeper sense. When she was not interacting with the world, she did not become dormant, but persisted and was active' (Kelly *et al.*, 2009, p. 318). Braude's more recent work (1995) corroborates the view that alters can be co-conscious. He points to the struggle of different alters for executive control of the body and the fact that alters 'might intervene in the lives of others [that is, other alters], intentionally interfering with their interests and activities, or at least playing mischief on them' (*ibid.*, p. 68). It thus appears that alters can not only be concurrently conscious, but that they can also vie for dominance with each other.

Clearly, the evidence indicates that different alters entail — to paraphrase Coleman (2014) again — different co-conscious points of view annexed to *private* qualitative fields, these private qualitative fields being carved out by virtue of dissociation. In other words, *different alters are different subjects*. The connected subgraph of phenomenal contents associated with an alter (see Figure 1b again) represents its private qualitative or phenomenal field. Moreover, alters form within a single overarching psyche, so the process of their formation entails a decomposition of an original subject into multiple lower-level subjects.

I submit that *dissociation in cosmic consciousness is what leads to the formation of relative subjects*. Each relative subject is thus an alter of cosmic consciousness, its private qualitative field corresponding to a segment of the latter's self-excitatory 'medium'.

By virtue of corresponding to a segment of cosmic consciousness, each alter retains — as Shani (2015) posited — the intrinsic features of *sentience* and *core-subjectivity*. But the local pattern of dissociative phenomenal activity in its respective segment is what bestows an alter its specific character, its unique perspective. In other words, the primary sense of I-ness of all alters is that of cosmic consciousness itself; the very consciousness of the alters, as an ontological 'medium', *is* cosmic consciousness. But the particular phenomenal field of an alter, which defines its identity as a seemingly separate individual, is demarcated by a local dissociative process — analogous to DID — in the corresponding segment of the 'medium'. Naturally, because alters are fully grounded in cosmic consciousness, it is incoherent to say that

they become separated from it; only an illusion of separation arises as a particular phenomenal content in the alter's dissociated qualitative field.

The key to my argument is the notion that dissociation can demarcate and carve out a private phenomenal field. This way, alters must become blind to all phenomenality taking place outside their respective field, which then explains why I cannot read your thoughts. And indeed, there is strong empirical evidence for the *literally blinding* power of dissociation: in 2015, doctors reported on the case of a German woman who exhibited a variety of alters (Strasburger and Waldvogel, 2015). Peculiarly, some of her alters claimed to be blind while others could see normally. Through EEGs, the doctors were able to ascertain that the brain activity normally associated with sight wasn't present while a blind alter was in control of the woman's body, even though her eyes were open. When a sighted alter assumed control, the usual brain activity returned. Clearly thus — if nothing else, for sheer empirical reasons — dissociation is a sufficiently powerful potential solution to the decombination problem.

10. At What Level Does Cosmic Dissociation Occur?

The challenge we must now address is the so-called 'boundary problem for experiencing subjects' (Rosenberg, 2004, pp. 77–90): what measurable structures in nature correspond to — that is, are the revealed appearance of — alters of cosmic consciousness? As we have seen, Shani (2015) posits that elementary particles are akin to micro-level alters, which in turn come together to compose higher-level relative subjects. However, as already mentioned, I believe this to be an unnecessarily convoluted notion. Instead, I submit that *cosmic dissociation happens precisely at the level of living beings with unitary consciousness, such as you and me*. You and I are alters of cosmic consciousness.

There are several arguments for this. The first has already been hinted at: given that we ordinarily experience an integrated phenomenal field, there is no direct reason to conjecture that this field is a composite of lower-level constituents.

Secondly, we have seen that von Neumann's reasoning regarding quantum measurement (1996) implies that the entire inanimate universe must be one unfathomable 'von Neumann chain' — that is, an entangled indivisible whole. As such, it is arbitrary — physically

speaking — to carve out any segment of the inanimate universe and posit it to be the revealed appearance of an alter. Von Neumann did, however, exclude conscious living organisms from the embrace of von Neumann chains, since at least *we*, conscious human beings, clearly *can* perform quantum measurements. On this basis, only conscious living organisms can correspond to alters of cosmic consciousness, not elementary particles or any other subset of the inanimate universe.

Thirdly, as observed by Mathews, ‘the individuation of [inanimate] objects... is not consistently objectively determined... many of our individuations — of rocks and mountains, for instance — have basically nominal status’ (2011, p. 144). Take what we call a ‘car’: though based on structural and functional reasoning that helps the business of transportation, its delineation is ultimately arbitrary. If one argues that, say, the spark plugs are integral to the car because without them the car cannot function, by the same token one would also have to include the fuel that makes its engine run, the environment air that allows combustion and cools the engine, the road gripped by the tires, the ground that sustains the road, the gravity that enables grip, and so on. The decision of where to stop is motivated by convenience. An analogous rationale applies to whether we distinguish the handle from the mug, the hood from the jacket, the river from the ocean, etc. This relative arbitrariness in the way we delineate their boundaries renders inanimate objects problematic candidates for the revealed appearance of alters of cosmic consciousness. After all, in Mathews’ words, ‘the boundaries between subjects are not nominal. The individuation of subjects, or centres of subjectivity, is objectively determined’ (*ibid.*).

Mathews is giving us an important clue here. Indeed, the boundaries of our own body are not nominal. Our ability to *perceive* ends at the surface of the body: our skin, retinas, eardrums, tongue, and the mucous lining of our nose. We cannot perceive photons hitting a wall or air pressure oscillations bouncing off a window, but we *can* perceive those impinging on our retinas and eardrums, respectively. Moreover, our ability to act through direct phenomenal intention also ends at the surface of the body: we can move our arms and legs simply by consciously *intending* to move them. However, we cannot do the same with tables and chairs. Clearly, thus, the delineation of our body is an *empirical fact*. I cannot just decide that the chair I am sitting on is integral to my body, in the way I *can* decide that the handle is integral to the mug. Neither can I decide that a patch of my skin is not integral to my body, in the way I *can* decide that the hood is not integral to the jacket. The criterion here is not merely a functional or

structural one, but the range of phenomenality — sensory perception, intention — intrinsically associated with our body. Based on this objective criterion, there is no freedom to move boundaries at will.

What these considerations suggest is clear: *the physical boundary of the body is the revealed appearance of the dissociative boundary of our phenomenal field*. And in so far as we can assume that all living organisms have phenomenal inner life in some way akin to our own, the conclusion can be generalized: *living organisms are the revealed appearance of alters of universal consciousness*; they are carved out of their context by virtue of cosmic dissociation.

But *can* we assume that all living creatures have phenomenal inner life? I believe we can: in so far as it resembles our own, the extrinsic behaviour of *all* metabolizing organisms is suggestive of their having dissociated phenomenal fields analogous to ours in some sense. This is obvious enough for cats and dogs, but what about plants and single-celled organisms such as amoebae? Well, consider this: ‘many types of amoeba construct glassy shells by picking up sand grains from the mud in which they live. The typical *Diffflugia* shell, for example, is shaped like a vase, and has a remarkable symmetry’ (Ford, 2010, p. 26). As for plants, many recent studies have reported on their surprisingly sophisticated behaviour, leading even to a proposal for a new field of scientific enquiry boldly called ‘plant neurobiology’ (Brenner *et al.*, 2006). Clearly, thus, even plants and single-celled organisms exhibit extrinsic behaviour somewhat analogous to our own, further suggesting that they, too, have dissociated phenomenal fields. Of course, the same cannot be said of any inanimate object or phenomenon (those that have been engineered by humans to merely *simulate* the behaviour of living beings, such as robots, natural language interfaces, etc., naturally don’t count).

Finally, we have good empirical reasons to believe that normal metabolism is essential for the maintenance of our dissociated phenomenal fields, for when it slows down or stops the dissociative boundary seems to become phenomenally porous (Kastrup, 2017a). So metabolism — the shared and differentiating characteristic of all living organisms — seems, again, to be the revealed appearance of alters of cosmic consciousness. The unique features of metabolism — think of DNA, morphogenesis, transcription, protein folding, mitosis, etc. — unify all life into a unique, clearly distinct natural category, despite the widely different forms organisms can take. This category provides the unambiguously demarcated ‘something in nature’ that Rosenberg was looking for (2004, pp. 77–90).

In conclusion, I posit that cosmic dissociation happens precisely at the level of individual living organisms. Each living organism is an alter of cosmic consciousness.

11. Reducing the Revealed to the Concealed Order

Notice that the revealed side of nature relative to any given alter consists of images on the screen of the alter's conscious *perception*. Therefore, if one can reduce perceptions at the level of individual alters to non-perceptual phenomenal contents at the level of cosmic consciousness as a whole, one will have reduced nature's revealed to its concealed order.

Before we address this challenge, however, we need some further background on dissociation. By definition, phenomenal contents inside an alter cannot evoke phenomenal contents outside the alter, and vice versa. But they can still *influence* each other. Indeed, phenomenal impingement across a dissociative boundary is empirically known. John Lynch and Christopher Kilmartin (2013, p. 100), for instance, report that dissociated feelings can dramatically affect thoughts and corresponding behaviours, whereas David Eagleman (2011, pp. 20–54) shows that dissociated expectations routinely mould our perceptions. Indeed, the entire clinical field of depth psychology is based on the notion that dissociated phenomenal contents in deeper layers of the psyche continuously impinge on the executive ego (Kelly *et al.*, 2009, pp. 301–34). We can visualize this as in Figure 2a, wherein the partial overlap of adjacent vertices internal and external to an alter represents impingement across its dissociative boundary.

Figure 2b illustrates the same thing according to a simplified representation unrelated to graph theory: the broader psychic space is represented as a white circle, with an alter represented as a grey circle within it. These circles are no longer graph vertices but represent sets of phenomenal contents. The dashed arrows represent the impingement of external and internal phenomenal contents — not explicitly shown — on each other, across the alter's boundary. For the avoidance of doubt, notice that these dashed arrows no longer represent cognitive associations. I shall use this simplified representation henceforth. Moreover, for simplicity's sake, from now on I shall also refer to all non-perceptual phenomenal contents simply as 'thoughts'.




Figure 2. Phenomenal contents impinging on the dissociative boundary of an alter, illustrated in two different but equivalent ways, (a) and (b).

I submit that, before its first alter ever formed, the only phenomenal contents of cosmic consciousness were thoughts. There were no perceptions. The formation of the first alter then demarcated a boundary separating phenomenal contents *within* the alter from those *outside* the alter. This newly formed boundary is what enabled perceptions to arise relative to an alter: the thoughts surrounding the alter *impinged* on its dissociative boundary from the outside. And since phenomenal contents are particular patterns of self-excitation of cosmic consciousness, this impingement can be regarded as an *interference pattern* between excitations within and outside the dissociative boundary, respectively (see Figure 2a again). What we call perception, or the revealed side of nature, is the alter's experience of this interference pattern (*cf.* Kastrup, 2017c). It follows that the revealed side of nature can be grounded in its concealed side: the former arises from excitatory interference between dissociated but mutually impinging *thoughts*. Indeed, I submit that the formation of dissociative boundaries is what partitioned the cosmos into revealed and concealed sides. See Figure 3.

The thoughts of an alter can also impinge on its dissociative boundary from the inside and thereby influence the surrounding phenomenal activity of cosmic consciousness (not shown in Figure 3). This corresponds to the effects on the world of the presence and actions of a living organism within it.




Figure 3. Thoughts in cosmic consciousness cause perceptions in an alter.

The revealed or extrinsic appearance of an alter’s boundary is an organism’s sense organs. In our case, these are our skin, eyes, ears, nose, and tongue. Therefore, even if the outside stimulation is very faint and subtle, evolution has had billions of years to optimize the sensitivity of our sense organs — our alter’s boundary — to pick up on these faint signals.

But how can a mere dissociative boundary give rise to a *qualitatively different* category of experience? If you recall, this is a question I raised earlier, motivated by the fact that perceptions *feel* undoubtedly very distinct from thoughts.

To answer it, let us first consider Donald Hoffman’s interface theory of perception (2009): it asserts that evolution emphasizes perceptual qualities conducive to *fitness*, not to truth. In other words, we have evolved to perceive not the phenomenal contents that are really out there — that is, outside our alter — but just a phenomenal *representation* thereof that helps us survive and reproduce. Hoffman uses the analogy of a computer desktop: although a computer file is represented in it as, for instance, a blue rectangle, this does not mean that the file itself has the qualities of being blue and rectangular. As a matter of fact, the actual file does not have those qualities at all: it is a pattern of open and closed microscopic switches in a silicon chip. In an analogous way, my hypothesis is that the qualities we experience on the screen of perception — colours, sounds, flavours, textures, etc. — are *not* the qualities experienced by the segment of cosmic consciousness that surrounds our alter, but their ‘desktop representation’ instead. Our perceptions do not feel like the thoughts of cosmic

consciousness because a *qualitative transition* between these two experiential categories has helped our ancestors survive and reproduce.

The work of Friston, Sengupta and Auletta (2014) has similar implications but, significantly, is derived from an entirely different line of reasoning. Their results are based on abstract mathematical considerations and, therefore, can in principle be leveraged under any ontology. They show that whenever a Markov blanket (Pearl, 1988) defines the boundary of an individual organism, internal states of the organism will evolve to optimize for two conflicting goals: (a) to reflect external states of the world beyond the Markov blanket; and (b) to minimize their own entropy or dispersion. Goal (a) is about allowing the organism to know what is going on in the world outside, so it can take suitable actions to survive in that world. Goal (b) is about preventing the organism from losing its internal structural and dynamical integrity because of the second law of thermodynamics. In our case, the dissociative boundary of an alter is the Markov blanket, whose revealed appearance is our skin and other sense organs.

The key insight of Friston, Sengupta and Auletta can be paraphrased as follows: a hypothetical organism with perfect perception — that is, able to perfectly *mirror* the phenomenal states of the surrounding external world in its internal states — would not have an upper bound on its own internal entropy, which would then increase indefinitely. Such an organism would dissolve into an entropic soup. To survive, organisms must, instead, use their internal states to actively *represent* relevant states of the outside world in a *compressed, coded form*, so to know as much as possible about their environment while remaining within entropic constraints compatible with their structural and dynamical integrity. This way, my hypothesis is that the qualities of perception experienced by an alter are just compressed, coded representations of how surrounding thoughts of cosmic consciousness are experienced from the concealed perspective. As such, while there must be a *correspondence* between perception and surrounding thoughts, the respective experiential qualities do not need to be the same. In fact, they will be very different if it helps organisms resist entropy. Our perceptions do not feel like thoughts because they are coded representations thereof.

12. Explaining the Correlations between Brain Function and Inner Experience

A principal argument for the mainstream physicalist position that the material brain somehow constitutes or generates consciousness is the empirically undeniable correlation between measurable brain function and inner experience (e.g. Koch, 2004). The way the idealist ontology proposed here accommodates this fact was already implicit in the previous section: *a metabolizing body — which includes a functioning brain — is simply the revealed appearance of the dissociated phenomenal field of an alter*. The former correlates with the latter simply because the former is what the latter *looks like* from across a dissociative boundary. Indeed, this can be empirically substantiated in a rather direct manner.

In a 2014 study of dissociation (Schlumpf *et al.*), doctors performed functional brain scans on both DID patients and actors simulating DID. The scans of the actual patients displayed clear differences when compared to those of the actors, showing that *dissociation has an identifiable extrinsic appearance*. In other words, there is something rather particular that dissociative processes *look like*. This further substantiates the notion that living organisms such as you and me are the revealed appearance of cosmic-level dissociative processes. After all, we now know empirically that dissociation is identifiable when observed from across the dissociative boundary. Metabolizing bodies are to dissociation in cosmic consciousness as certain patterns of brain activity are to DID patients.

Let me elaborate further on this important point. For any given alter *A1* of cosmic consciousness, it is the phenomenal contents surrounding *A1* that cause its perceptions of the world around it. Dissociated phenomenal contents corresponding to another alter *A2* can be part of the phenomenal environment surrounding *A1*. As such, the inner experiences of *A2* can also indirectly stimulate *A1*'s boundary — by impinging on their shared phenomenal environment — and thereby cause *A1*'s perceptions of *A2*. This is what gives *A1* access to the revealed appearance of the inner experiences of *A2* in the form of *A2*'s metabolizing body. See Figure 4. And since *A2*'s brain is integral to its body, it follows that *A2*'s inner experiences cause the perception by *A1* of the activity in *A2*'s brain. This causal link explains the correlations between inner experience and corresponding patterns of brain activity.




Figure 4. A metabolizing body is the revealed appearance of an alter's dissociated phenomenal field.

In essence, the claim here is that there is nothing to a metabolizing body but the revealed side — the extrinsic appearance — of the corresponding alter's inner experiences. Yet, one may object to this by arguing that many parts of the body seem entirely unrelated to inner experience: whereas certain patterns of brain activity correlate with subjective reports of experience, a lot seems to go on in the brain that subjects have no introspective access to (Westen, 1999; Hassin, Uleman and Bargh, 2005; Dijksterhuis and Nordgren, 2006; Augusto, 2010; Hassin, 2013). Moreover, what kind of inner experience does, say, liver function correspond to? What about big-toe function?

The answer to this objection is precise and compelling, but elaborate and specialized enough to have required its own paper (Kastrup, 2017b). Here, I shall simply remind the reader that a subject's lack of *metacognitive access* to an experience precludes reporting of the experience to self or others, but does not imply *absence* of the experience from the subject's qualitative field. With the emergence of no-report paradigms in neuroscience (Vandenbroucke *et al.*, 2014; Tsuchiya *et al.*, 2015), we now know that much is experienced that cannot be reported even to self, for subjects are often not aware *that* they have certain experiences. Moreover, as mentioned earlier, there are normal internal dissociations in the human psyche — the foundational claim of depth psychology — that render much of its phenomenal contents inaccessible to the reporting ego (Kelly *et al.*, 2009, pp. 301–34). So the hypothesis I am positing here is not defeated by the objection: all bodily metabolism — yes, even liver and toe function —

can still correspond to concealed phenomenal contents, even though these contents may not be introspectively accessible.

13. Explaining Our Shared World

The final explanatory burden that needs to be addressed is the undeniable empirical fact that we all inhabit seemingly the same environment, and that the laws that govern the dynamics of this environment operate independently of our personal volition. After all, if the world is imagined — as implied by idealism — how come we are all imagining seemingly the same autonomous world?

Notice that the existence of a phenomenal environment wherein all metabolizing organisms are immersed — a shared world — is a direct implication of the argument already developed. To bring this out, we simply need to extend Figure 3 to multiple alters, as illustrated in Figure 5. All alters are immersed, like islands of a single ocean, in the thoughts that constitute the concealed side of the inanimate cosmos. These thoughts surround all alters and cause their mutually-consistent perceptions by impinging on their respective dissociative boundaries. And since the volition of an alter is a phenomenal content also dissociated from the rest of cosmic consciousness, it follows that alters cannot change the laws of nature. From the dissociated perspective of alters, the world is thus autonomous.




Figure 5. Alters are immersed in a common phenomenal environment.

14. Conclusions

I have elaborated on an idealist ontology that can be summarized as follows. There is only cosmic consciousness. We, as well as all other living organisms, are but dissociated alters of cosmic consciousness, surrounded by its thoughts. The inanimate world we see around us is the revealed appearance of these thoughts. The living organisms we share the world with are the revealed appearances of other dissociated alters. This idealist ontology makes sense of reality in a more parsimonious and empirically rigorous manner than mainstream physicalism, bottom-up panpsychism, and cosmopsychism. It also offers more explanatory power than these three alternatives, in that it does not fall prey to the hard problem of consciousness, the combination problem, or the decombination problem, respectively.

References

- American Psychiatric Association (2013) *Diagnostic and Statistical Manual of Mental Disorders*, 5th ed., Washington, DC: American Psychiatric Publishing.
- Augusto, L.M. (2010) Unconscious knowledge: A survey, *Advances in Cognitive Psychology*, **6**, pp. 116–141.
- Black, D.W. & Grant, J.E. (2014) *The Essential Companion to the Diagnostic and Statistical Manual of Mental Disorders*, 5th ed., Washington, DC: American Psychiatric Publishing.
- Braude, S. (1995) *First Person Plural: Multiple Personality and the Philosophy of Mind*, New York: Routledge.
- Brenner, E.D., et al. (2006) Plant neurobiology: An integrated view of plant signaling, *Trends in Plant Science*, **11** (8), pp. 413–419.
- Carruthers, P. & Schechter, E. (2006) Can panpsychism bridge the explanatory gap?, *Journal of Consciousness Studies*, **13** (10–11), pp. 32–39.
- Chalmers, D.J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*, New York: Oxford University Press.
- Chalmers, D.J. (2003) Consciousness and its place in nature, in Stich, S. & Warfield, F. (eds.) *Blackwell Guide to the Philosophy of Mind*, Malden, MA: Blackwell.
- Chalmers, D.J. (2016a) The combination problem for panpsychism, in Brüntrup, G. & Jaskolla, L. (eds.) *Panpsychism*, Oxford: Oxford University Press.
- Chalmers, D.J. (2016b) Panpsychism and panprotopsychism, in Brüntrup, G. & Jaskolla, L. (eds.) *Panpsychism*, Oxford: Oxford University Press.
- Coleman, S. (2014) The real combination problem: Panpsychism, micro-subjects, and emergence, *Erkenntnis*, **79** (1), pp. 19–44.
- Dijksterhuis, A. & Nordgren, L.F. (2006) A theory of unconscious thought, *Perspectives on Psychological Science*, **1** (2), pp. 95–109.
- Eagleman, D. (2011) *Incognito: The Secret Lives of the Brain*, New York: Canongate.
- Eddington, A.S. (1928) *The Nature of the Physical World*, New York: The MacMillan Company.

- Ford, B.J. (2010) The secret power of the single cell, *New Scientist*, **206** (2757), pp. 26–27.
- Friston, K., Sengupta, B. & Auletta, G. (2014) Cognitive dynamics: From attractors to active inference, *Proceedings of the IEEE*, **102** (4), pp. 427–445.
- Goff, P. (2006) Experiences don't sum, *Journal of Consciousness Studies*, **13** (10–11), pp. 53–61.
- Goff, P. (2009) Why panpsychism doesn't help us explain consciousness, *Dialectica*, **63** (3), pp. 289–311.
- Hassin, R.R. (2013) Yes it can: On the functional abilities of the human unconscious, *Perspectives on Psychological Science*, **8** (2), pp. 195–207.
- Hassin, R.R., Uleman, J.S. & Bargh, J.A. (eds.) (2005) *The New Unconscious*, New York: Oxford University Press.
- Hoffman, D.D. (2009) The interface theory of perception: Natural selection drives true perception to swift extinction, in Dickinson, S., et al. (eds.) *Object Categorization: Computer and Human Vision Perspectives*, Cambridge: Cambridge University Press.
- Horgan, T. & Potrč, M. (2000) Blobjectivism and indirect correspondence, *Facta Philosophica*, **2** (2), pp. 249–270.
- Jaskolla, L.J. & Buck, A.J. (2012) Does panexperiential holism solve the combination problem?, *Journal of Consciousness Studies*, **19** (9–10), pp. 190–199.
- Kastrup, B. (2017a) Self-transcendence correlates with brain function impairment, *Journal of Cognition and Neuroethics*, **4** (3), pp. 33–42.
- Kastrup, B. (2017b) There is an 'unconscious,' but it may well be conscious, *Europe's Journal of Psychology*, **13** (3), pp. 559–572.
- Kastrup, B. (2017c) Making sense of the mental universe, *Philosophy and Cosmology*, **19**, pp. 33–49.
- Kelly, E.F., et al. (2009) *Irreducible Mind: Toward a Psychology for the 21st Century*, Lanham, MD: Rowman & Littlefield.
- Koch, C. (2004) *The Quest for Consciousness: A Neurobiological Approach*, Englewood, CO: Roberts & Company Publishers.
- Levine, J. (1983) Materialism and qualia: The explanatory gap, *Pacific Philosophical Quarterly*, **64**, pp. 354–361.
- Lynch, J. & Kilmartin, C. (2013) *Overcoming Masculine Depression: The Pain Behind the Mask*, New York: Routledge.
- Mathews, F. (2011) Panpsychism as paradigm, in Blamauer, M. (ed.) *The Mental as Fundamental*, Frankfurt: Ontos Verlag.
- Nagasawa, Y. & Wager, K. (2016) Panpsychism and priority cosmopsychism, in Brüntrup, G. & Jaskolla, L. (eds.) *Panpsychism*, Oxford: Oxford University Press.
- Neumann, J. von (1996) *Mathematical Foundations of Quantum Mechanics*, Princeton, NJ: Princeton University Press.
- Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, San Francisco, CA: Morgan Kaufmann.
- Rosenberg, G. (2004) *A Place for Consciousness*, New York: Oxford University Press.
- Russell, B. (2007) *The Analysis of Matter*, Nottingham: Spokesman Books.
- Schaffer, J. (2010) Monism: The priority of the whole, *Philosophical Review*, **119** (1), pp. 31–76.

- Schlumpf, Y., *et al.* (2014) Dissociative part-dependent resting-state activity in Dissociative Identity Disorder: A controlled fMRI perfusion study, *PLoS ONE*, **9**, doi:10.1371/journal.pone.0098795.
- Seager, W. (2010) Panpsychism, aggregation and combinatorial infusion, *Mind and Matter*, **8** (2), pp. 167–184.
- Shani, I. (2015) Cosmopsychism: A holistic approach to the metaphysics of experience, *Philosophical Papers*, **44** (3), pp. 389–437.
- Strasburger, H. & Waldvogel, B. (2015) Sight and blindness in the same person: Gating in the visual system, *PsyCh Journal*, **4** (4), pp. 178–185.
- Strawson, G. (2006) *Consciousness and Its Place in Nature*, Exeter: Imprint Academic.
- Tsuchiya, N., *et al.* (2015) No-report paradigms: Extracting the true neural correlates of consciousness, *Trends in Cognitive Sciences*, **19** (12), pp. 757–770.
- Vandenbroucke, A., *et al.* (2014) Seeing without knowing: Neural signatures of perceptual inference in the absence of report, *Journal of Cognitive Neuroscience*, **26** (5), pp. 955–969.
- Westen, D. (1999) The scientific status of unconscious processes: Is Freud really dead?, *Journal of the American Psychoanalytic Association*, **47** (4), pp. 1061–1106.

Paper received July 2017; revised February 2018.