Locating Consciousness: Why Experience Can't Be Objectified

Thomas W. Clark

Correspondence: twclark@brandeis.edu

Abstract: *The world appears to conscious creatures in terms of experienced sensory qualities, but science doesn't find sensory experience in that world, only physical objects and properties. I argue that the failure to locate consciousness in the world is a function of our necessarily* representational relation *to reality as knowers: we won't discover the terms in which reality is represented by us in the world as it appears in those terms. Physicalists who are realists about consciousness generally assume its objectivity: experience is something identical with physical processes or properties, perhaps the intrinsic nature of the physical, or perhaps some micro-physical, neural, or emergent property. I argue that this assumption wrongly reifies consciousness; it expects to find qualitative representational content – qualia – in the physical world as characterized using such content. Instead, we should grant that conscious experience constitutes a mind-dependent, subjective,* representational *reality for cognitive systems such as ourselves, and that the physical world described by science is a* represented *objective reality. The former, since it exists only for conscious subjects, won't be found as an entity in the latter. I suggest that naturalistic approaches to explaining consciousness should acknowledge the representational relation and the non-objectivity of experience, and be constrained by evidence that consciousness accompanies certain sorts of behavior-controlling representational functions carried out by complex, physically-instantiated mind-systems. I evaluate a variety of current hypotheses about consciousness on that basis, and suggest that a mature science of representation may eventually help explain why, perhaps as a matter of representational necessity, experience arises as a natural but not objectively discoverable phenomenon.*

1. The privacy constraint on consciousness

On a naturalistic view of ourselves, we are entirely physical beings who are also conscious, but thus far there is no consensus on the nature of consciousness. The central difficulty is that the defining feature of conscious experience – the subjective, qualitative 'what it's like' or phenomenal character of tasting a mango, seeing a red rose, or dreaming about a blue lake – is not available to intersubjective observation or measurement (Gamez, 2014). If it were, there would be no problem of consciousness, nor of other minds, for instance whether fish feel pain. Pain would be out there in public, so we'd know that they either do or don't suffer when hooked. But pain isn't public, unlike whatever its neural or otherwise physical correlates might be, whether in fish, fowl, or us. We don't and won't see pain when peering into the brain, at whatever level of magnification. And so it is with all experiences: they are only available to, only exist for, individual conscious subjects. Unlike the subjects' brains, they aren't observables.

Despite its subjective nature, many approaches to explaining consciousness hypothesize that experience is an objective, physically-embodied phenomenon; optimistic physicalists suppose that consciousness will eventually find its place in the material world as described by science.

Theories aiming to objectify consciousness range from reductive identity theories that equate experiences with physically-realized states or functions; non-reductive, property dualist theories in which consciousness supervenes on material or functional states; panpsychist theories in which phenomenology is a fundamental property of matter; theories involving quantum or magnetic fields as instantiators of consciousness; and radically enactivist theories that identify conscious states with ordinary physical objects (Manzotti, 2011).

Although we can't conclusively discount these possibilities, I'd suggest we not hold our collective breath. The existential privacy of an experience, its subjectivity, that *it exists only for the mind undergoing it*, isn't likely going away. We can call this the privacy constraint on consciousness. We can then ask why conscious experience isn't objectifiable even though as conscious creatures we are physically objective. The answer I propose has to do with what I'll call the representational relation: the world is only *known* by cognitive systems, including ourselves, using content-bearing representations. Conscious experience is arguably a species of representational content – *qualitative, phenomenal* content – and the world appears to each of us as a conscious subject in terms of that content. The reason we don't find conscious experience in the world, the reason we can't objectify it, is because as a rule we don't and won't find representational content in the world as modeled by it. We only find the physical objects and phenomena characterized *in terms of* such content, including the physically-instantiated content vehicles. Theories that suppose we can objectify consciousness, putting it in the public domain, are thus barking up the wrong tree.

Still, the puzzle remains of why, as current evidence strongly suggests, only certain sorts of physically and functionally specified cognitive systems, running certain sorts of representational processes in service to behavior control and system integrity (T. Clark, 2005, 52-5), end up hosting conscious experience. We can perhaps narrow down the explanatory possibilities in light of the realization that we don't, and won't, discover consciousness to be among the occupants of the world as it appears to us, either in experience or as described by science.

The plan of the paper is as follows: Section 2, next up, argues that as knowers we are necessarily in a representational relation to reality, and that conscious experience, because it arises in conjunction with sensory-perceptual content-bearing representational processes in the brain, can itself be understood as bearing representational content. Section 3 argues, against illusionism, that experiential content is genuinely qualitative in character, such that conscious phenomenology is no illusion, but something in need of explanation by a theory of consciousness. That explanation must also account for the privacy constraint: the fact that qualitative, phenomenal content – qualia – is not an intersubjectively available feature of the world, but available only to the conscious subject. Section 4 describes the intersubjective representational project of objectification that locates phenomena in the world, a project exemplified by science but prefigured in shared folk descriptions of physical objects. To be a physical phenomenon in space (spacetime in physics) is to be *represented as such* in the everyday terms of qualitative experience and the more specialized, abstract terms of scientific quantitative descriptions.

However, the fact remains that science can't seem to locate consciousness, and Section 5 suggests why not: as a general rule representational content won't *itself* be found in the world as represented in terms of such content, whether conceptual, propositional, quantitative or qualitative. Having proposed that the unlocatability and non-objectivity of phenomenal content is

an instance of this rule, I then proceed in Section 6 to evaluate some current theories of consciousness in light of the non-objectivity thesis, and highlight what I think are among the most promising approaches. In particular I favor Thomas Metzinger's well-developed 'self-model theory of subjectivity'. However, Section 7 takes issue with some of Metzinger's claims about consciousness and describes the advantages of adopting a non-interactionist phenomenal-physical parallelism. In the last section I conclude by suggesting that a mature science of representation might eventually do the trick in explaining consciousness.

2. The representational relation

It is a commonplace that as knowers we deploy various sorts of representations in negotiating our contact with the world. The world is represented by us conceptually and quantitatively in the 'manifest image' of ordinary human discourse and the 'scientific image' of physical theory (Sellars, 1962). That we are in a representational relation to reality seems an unavoidable condition of our being limited, situated creatures with particular perspectives on the world, whether individual or collective. This means that the world never appears to us undressed, so to speak, but always clothed in perspectivally conditioned models. Still, the models we humans deploy generally include a vague but plausible (and perhaps unavoidable) realist assumption: the world exists mind-independently and includes various mind-independent entities and processes, some of which appear to us as having spatio-temporal, physical properties as given in both science and everyday experience. Among those entities are composite, complex, and integrated systems that constitute minds – mind-systems – at least some of which, like ourselves, are conscious.

On the face of it, conscious experience seems to be a representational, informationally rich phenomenon that mediates our contact as individual subjects with the world. There's usually a non-coincidental and behaviorally crucial correspondence between our waking experience and what's the case in our immediate surroundings. This correspondence is underwritten by causal interaction with the environment via our information-gathering, behavior-guiding sensory modalities, the operation of which consciousness is closely associated. Most of the time we unreflectively take the world as given in experience to be the spatio-temporal manifold as it is in itself, directly presented to us. But we can infer, on the basis of dreams, hallucinations, and optical illusions that experience is a selective and fallible individual-level model of what's outside the head.

This shouldn't be construed as saying that we somehow see *experience* – the model – instead of the world; we shouldn't suppose we observe consciousness (T. Clark, 2005). We can avert our gaze and otherwise perceptually distance ourselves from physical objects, but cannot divorce ourselves from the experience in terms of which objects appear and disappear for us since, as subjects, we *consist* of experience. To be conscious is for us to subjectively constitute an experiential world-model (Revonsuo, 2015) – what Thomas Metzinger (2009) calls an 'ego-tunnel' – that is modulated and constrained by our direct, physical contact with the world itself (including the body) via our sensory-perceptual systems.

We have developing theories of such contact, most recently and notably Bayesian predictive coding (Seth et al., 2011; A. Clark, 2013): impinging stimuli activate sensory channels that inform the brain's current multi-modal world-model, helping to minimize mismatches between neural representations and the world in service to behavior control. The continually updated mappings and covariances between the world and brain – the neurally-realized representational

relation of sensory perception – allows for successful action and system maintenance, given the nature and needs of the organism (Kanwisher, 2001; Dehaene & Changeux, 2011; Sterling & Laughlin, 2015).

The information-bearing neural processes associated with conscious experience – call them *conscious* processes – can be identified by contrasting the neural networks active when performing tasks only possible when conscious (e.g., complex learning, planning, reporting) to those networks subserving behavior that can be handled unconsciously (e.g., habitual or automatized behaviors) (Baars, 1997; T. Clark, 2005, 52-55). Experience, since it closely correlates with conscious processes that carry information about the world, *itself* tracks the world, at least when we're awake and in perceptual contact with our body and environment. Consciousness can thus carry representational content as inherited from its neural correlates, but couched in qualities available only to the subject (the privacy constraint).[1]

3. The reality of phenomenal content

Despite the apparent reality of consciousness as constituting a qualitatively rendered world-model, one approach to solving the problem of how it fits *into* the world as thus modeled is to deny the existence of qualitative experience itself. Illusionists such as Keith Frankish (2016) say that phenomenal qualities – the 'introspectable qualitative properties, or "feels", which determine what it is like to undergo [experiences]' – simply don't exist. According to illusionist Daniel Dennett, what's illusory are the 'special subjective properties (typically called qualia) to which you (alone) have access'(2017, 365), properties that 'we are acquainted with more directly, when we are slightly less directly acquainted with their normal external causes' (2017, 360). These properties are thought to be

> …an *internal cause* that has the same properties as the *intentional objects* that normally cause your perceptual beliefs – except that these are private, subjective versions, somehow, of the public, objective properties of redness and so forth. (2017, 361, original emphasis)

Illusionists say there are no such private, subjective properties or causes – no qualia in the sense that Dennett says they are typically conceived; there only seem to be. Nevertheless, Dennett says (and I agree) that there exist representational contents that individuate our experiences:

> …if you want to talk about your own mental states, you *must* identify them by their content: 'Which idea? My idea of HORSE. Which sensation? My sensation of *white*'. (Dennett 2017, 367, original emphasis)

Thus, to use Dennett's (2016, 2017) example, when we experience an afterimage – a red stripe produced by looking at an image of a green- and white-striped American flag – there really exists a bit of representational content (*red*), but not any *instantiated property*. We won't find the afterimage anywhere – not in the brain, not out there in the world – so there's no instantiation of redness. But we all agree that something appears red, which attests to the reality of the representational content which participates in individuating our afterimage experience; the content *red* remains a perfectly real feature of our experience as we all report it.

[1] I won't address here the question of whether *all* conscious episodes, including all qualitative feels, carry representational content, but it seems clear enough that some do.

Although it may not involve subjective properties or entities as Dennett portrays them (internal causes that instantiate properties to which we alone have access), experiential content nevertheless presents itself as *qualitative*, as contrasted, for example, with the propositional content of mental states such as beliefs and desires. Dennett (2015) sketches an explanation of why, as he puts it, '…qualia seem so simple and ineffable,' why they seem 'atomic to introspection, unanalyzable simples' – the hallmark of qualities. It is, he says

> …clearly an effect of something like the resolution of our discernment machinery…It turns out that these 'qualia' are actually highly structured properties of neural representations. The explanation of this effect, needless to say, is ultimately in the medium of [neural] spike trains.

Whatever the explanation of this (real) effect turns out exactly to be, qualitative content in its simplest particulars is the basic 'what it's like' of having experiences: the non-decomposable, atomic quality of primary red, the root, not-further-specifiable quality of pain, whatever its type.

Frankish (2016) allows that experiences, construed functionally as 'the mental states that are the direct output of sensory systems', have representational content, but not *qualitative* content; the qualitative aspect of experience is, he says, illusory. But I would argue that for the experiencer it's *no illusion* that there are no further discriminable components of basic pain or primary red available to her; such non-decomposability is, again, the hallmark of qualitativeness. If something appears to be qualitative in experience, that appearance counts as qualitative (what else would it count as?) and thus needs explanation, for example as Dennett sketches above and as I will sketch in Section 6 when discussing Metzinger.

If qualitative content exists, then so does phenomenology and phenomenal consciousness. Phenomenology, simple or complex, is constituted by the qualities which individuate conscious experiences, and those qualities (often, perhaps not always, see note 1) exist as representational, information-bearing features of experience whether or not any object exists to which the content, e.g., *red*, is attributed to as a property. If we ask where *is* the red of Dennett's afterimage, the correct answer is: *nowhere*. But the red is still perfectly real as a bit of representational content that persists as long as the conscious episode persists.[2]

Looking at the (non-illusory) red apple over there, we naturally want to say the red is over there too; we certainly *represent* it to be over there. But what Dennett's afterimage suggests (and I will argue further below) is that the phenomenal content in terms of which we characterize physical objects in our experience *isn't anywhere*. Of course, unlike the afterimage, we commonsensically take the redness of the apple as its own mind-independent property. But as much as experienced colors pick out and characterize physical objects, and thus are their *represented* properties (properties we represent objects as having), colors and other sensorily delivered properties of objects are in the first instance unlocatable phenomenal contents. There is a real, locatable apple with objectively specifiable features (mass, chemical composition, etc.), but it appears to us in

[2] We might be tempted to say that the qualities given in conscious episodes are *qualities of experiences*, except for the fact that representational content is intentional, that is, usually of about or something other than the episode itself. (Exceptions are of course when we engage in meta-representational thought or discourse about representations themselves, as exampled by this paper.) In the case of the afterimage, the experienced red is of or about the (illusory) afterimage, not the experience itself, which in part *consists* of the content *red*.

conscious experience in terms of (real) phenomenology that isn't itself anywhere. This claim is consistent with, and indeed is a reformulation of, the privacy constraint on consciousness.

However we eventually account for experienced qualities, in what follows I'll use 'qualia' to refer to them – the original sense of the term as introduced by C.I. Lewis (1929); we need not follow Dennett in his qualia skepticism if, prior to adopting any metaphysical view about consciousness, we simply define them as the basic irreducible particulars of experiences. And I hope to have motivated *qualia realism* (Goff, 2016): for individual conscious subjects, the physical world appears in all its extended, massive, colorful glory in terms of real phenomenal content.

4. Objectification and the physical

The *prima facie* privacy of conscious experience is that my pain (as phenomenal content, not the associated neural goings-on or overt behavior) exists only for me; this contrasts with the public accessibility and observability of physical objects like my brain. For something to count as physical – for us to represent it as having objective, mind-independent existential status as a physical phenomenon – is for it to be actually or potentially intersubjectively accessible (observable, perceivable, measurable). Before science steps in, the physical objectivity of things like apples gets certified by intersubjective consensus on descriptions couched largely in the qualitative terms of everyday sensory experience: we generally agree about the apple's redness, roundness, surface texture, firmness, and perhaps smell and taste, all properties we commonsensically attribute to it, even though, as noted above, they are *represented* properties that participate as contents of our phenomenal world-model. The apple is concretely, tangibly present: it has physical substance – 'stuff' – as given in our experience. The apple thus achieves its objective, intersubjectively available status as what I'll call a *folk-physical* object. After science gets involved, however, the descriptions include quantitative parameters that make no mention of experienced qualities: the pH level of the apple's acidity, the relative concentration of certain esters, its rest mass, reflectance properties, etc. Eventually, the scientific description of the apple as a physical object might leave behind any qualitatively rendered, experienced component at all – it has become a *science-physical* object.

We can think of this transition, from descriptions of objects in terms of qualities (folk-physical) to entirely quantitative descriptions (science-physical), as a matter of increasing objectification. The phenomenal qualities in folk descriptions of objects can of course vary from individual to individual, depending on their sensory-perceptual capacities. Color-blind individuals may not be able to reliably discriminate a red from a green apple, and if you've lost your sense of smell, cider will taste merely sweet, lacking its characteristic appley-ness that would allow you to discriminate it from orange juice (texture and mouth feel aside). The descriptions afforded by science aren't prone to such variability since the objects they pick out are reliably identified as having quantifiable properties according to reproducible measurements that in principle anyone or any*thing* could carry out, including aliens and AIs. Because we ordinarily think of reality as having its own mind-independent nature, descriptions that leave behind experienced, mind-dependent qualities in favor of quantifiable characteristics are in that sense more objective.

We should not forget, however, that the maximally objective, quantified specification of the world – the spatio-temporal world as described by science – is still a representationally encoded *model*, not the unrepresented world itself. The unrepresented self-nature of reality, what we intuitively believe exists independently of our representations of it, is still at an epistemic

remove, even though we ourselves participate in that reality. At its best, what the scientific model can afford us is a predictively successful and explanatorily consistent *structurally isomorphic rendition of reality*, one couched in terms of physical parameters and constants (Ladyman and Ross, 2007). Unrepresented self-natures – Kantian things in themselves – are necessarily left out of the maximally objective picture of the world since the epistemic, representational interface between knowers and known is always in place.

5. The non-objectivity of representational content

Having maximally objectified the world, science seems to stumble when it comes to consciousness. There is simply nothing in the world as described by our most objective account of it that answers to what's diagnostic of consciousness: qualitative phenomenology. If there were, the problem of consciousness would not arise. However, if we grant the reality of phenomenal content, we naturally want to objectify it: locate it in the physical world as given to us in our best scientific models. That we can't affronts the explanatory ambition to unify, for instance under physicalism, all that we hold to be real.

To soften the blow to physicalism, we can note that it isn't just qualia that we don't find in the world they help represent for us. Although numbers are indispensable in quantifying reality, we don't expect to see them sitting in spacetime. Rather, we take them to be mind-dependent representational tools that we use in *descriptions* of objects, not objects themselves. We can count cats and black holes, but not the numbers we use to count them since numbers aren't anywhere to be counted. We can write down an equation, but what we see on the whiteboard is a concrete physical phenomenon, the vehicle of representation, not the mathematical content itself. It is sometimes claimed that the reality of math is not just a matter of its representational function, but that mathematical objects exist in a mind-independent Platonic realm. This seems to me a misguided attempt to reify representational content: we don't need to reify numbers in any sense to count them as representationally indispensable.

Likewise for concepts and propositions. There are millions of objects picked out by the concept CAR, but the concept won't be found on roadways or anywhere else. Nor will we find the proposition just proposed, only the physical and cultural circumstances that make it true: cars – what CAR refers to – exist as locatable physical objects, concepts do not. But that we so confidently and necessarily speak of concepts and propositions, given our representational proclivities, suggests that even if they aren't locatable, they participate in what is for us a collective *representational reality*. Although they don't appear in the physical world, we can't transcend concepts and quantities when constructing our world-models, folk or scientific (although they might change as our models improve). Likewise, following Thomas Metzinger's view (about which more below), I suggest we can't transcend the phenomenal world-model associated with our sensory-perceptual capacities. Experience too constitutes a representational reality, albeit personal, not collective, and, like concepts and propositions, we don't find experiential content in the world as experience or science represents it (the privacy constraint). We of course might find in the world what's *referred to or picked out* by representational episodes, what we might call the *referential* content (e.g., the cars we were just thinking about), but not the *episodic* content as it (momentarily) exists in conscious episodes such as perceptions, thoughts, dreams or hallucinations.

The forgoing suggests what might be a general rule about representational content, of which the privacy constraint on consciousness is one instance: the terms in which we represent

(characterize, describe, grasp) the world will not be found among the phenomena of the world they participate in representing. Put another way, content, although real, cannot be objectified, such that we will find it alongside the physical objects and processes (whether folk-physical or science-physical) that appear to us in terms of that content. The constituents of our *representational* reality – concepts, propositions, numbers and, finally, qualia – don't and won't appear as objects or substances, whether physical or non-physical, in the reality they represent as being objectively the case, our *represented* reality. Physical phenomena, including the vehicles of representation (printed words, diagrams, pictures, neural nets) are represented by us as being spatial, concrete, extended, etc. using terms (qualia, concepts, numbers) that aren't themselves spatial, concrete, or extended.

Against the claim that content can't be objectified, it might be objected that we *can* observe representational content in information-bearing systems. We could perhaps specify what states in what processors running a face recognition program code for *Tony Blair face*, content that then gets used in labeling him in social media images. Wouldn't that specification just *be* the content itself? But it's the successful *performance of the system* as judged by its concrete output (image, text and auditory tokens) that ends up validating the ascription of content; nothing in the processing looks like him such that you can see the content in the system. All you see is processing and processors – the content-bearing vehicles, their organization, and the resulting behavior of the system as it labels him in images. The system's episodic content (that which exists when the system is up and running) isn't causally produced by the system's representational vehicles as a further physical effect – there's no 'second transduction', as Dennett (1998) might put it – so the content doesn't exist as an intersubjectively accessible observable.

On the plausible (to me) assumption that neural activity sometimes carries representational content, whether unconscious or conscious, the same point applies: you won't see content percolating in my brain, e.g., my thought that the apple is on the table or my experience of pain; you'll just see neurons doing their thing. Neurally-instantiated content, whether at the sub-personal level (e.g., single channel sensory information) or personal (e.g., beliefs, desires) is *proprietary* to the human organism, not publicly accessible: it exists as a function of the organism's self-regulation and behavior control processes. It is thus system-dependent, even though the specific content will often reflect the organism's ongoing interaction with its environment. The content, being representational, will of course often refer to or pick out publicly accessible states of affairs. I necessarily think of things *in terms of* episodic content, e.g., my thinking of (visualizing, imagining) the magnolias now in bloom in Boston. But you can only see (access) the magnolias, not the content of my thoughts concerning them.

Cross-species commonality in brain function and structure, for instance among primates and other mammals, strongly suggests that hosting *conscious* content – phenomenal experience – is not just a human prerogative (Feinberg & Mallatt, 2018; Ginsberg & Jablonka, 2019). But knowing the point at which a system, natural or artificial, ends up with a full-blown self-in-the-world reality-model awaits a credible theory of consciousness, and consciousness may not be an all or nothing affair.[3] Still, since the qualitative content of conscious episodes isn't an observable, the inference to similarity in content (conscious or not) across species and artificial systems is just that, an inference based on their observable features and behavior (Dennett,

[3] On consciousness as a 'graded phenomenon' see Metzinger (2003), 135.

1987). One has to *be* a certain sort of system, the nature of which is under investigation (Dehaene & Changeux 2011; Prinz, 2012; A. Clark, 2013), for the content to become, for it, the private representational reality of conscious experience.

6. Evaluating hypotheses on consciousness

If conscious experience constitutes a subjective representational reality, what in the *represented, objective* reality afforded by science and philosophy (the philo-scientific image, we might call it) could explain the existence of experience for mind-systems? The claim that I've just defended – that phenomenal content is real but can't be objectified – might constrain our approaches to explaining consciousness. Other constraints mentioned above include the assumption of a mind-independent reality that includes mind-systems like ourselves; the fact that consciousness seems to be a mind-*dependent* phenomenon; and the fact that minds, as far as we know, are composite, complex systems which model reality in service to self-maintaining behavior. In what follows, I will bring these considerations to bear in assessing the viability of hypotheses about consciousness. My selection of hypotheses to evaluate is necessarily incomplete given the vast landscape of consciousness studies, and the evaluations themselves will be cursory but I hope suggest promising avenues for exploration.[4]

6.1 Non-starters

If we don't expect to find consciousness in the (represented) physical world, this rules out hypotheses which hold that experience is causally generated by physical states of affairs. Anything caused by physical goings-on will itself be physical – a potentially observable and thus objective phenomenon. As noted above, there is no 'second transduction' in the brain that produces content, including the phenomenal content of conscious experience, as a further physical effect of neural processing. Consciousness isn't generated as a measurable output of its correlates, whatever those turn out precisely to be, so there will be no *causal* explanation of qualia forthcoming (Oakley & Halligan, 2017). Epiphenomenalists who suppose that qualia are somehow caused by physical states of affairs, but then play no causal role in behavior (Robinson, 2010) face this problem: there is no evidence for any causal production or generation of mental phenomena by the physical, so the epiphenomenalist worry seems misplaced until it's established that consciousness is indeed objective – something that exists on the same causal playing field as the brain.

Likewise, non-reductive physicalists who propose that consciousness somehow emerges from or supervenes on neural processing (or any other presumptive physical base) as an objective but as yet unobservable property or feature of such processing must specify what mechanisms and transitions are involved. Such property dualism has it that consciousness can be individuated as a distinct mental aspect of the physical states of affairs with which it is associated, something that doesn't reduce to them but nevertheless exists alongside them in the objective world. In this way, it's possible to maintain that consciousness plays its own behavior-controlling role beyond what non-mental properties accomplish (which is how it might commonsensically seem). But then the vexed problem of phenomenal causation arises: how exactly does qualitative experience – something invisible to science – causally contribute to physically-mediated behavior? In any case, if consciousness can't be objectified, there is no objective mental property which could

[4] For compendia on hypotheses about consciousness, see Blackmore (2006), Van Gulick (2014), and Bourget and Chalmers (2018).

play a behavior-controlling role. And indeed, no such property has been discovered; all we have in the objective (represented) world of scientific descriptions (the best model going) are the physical entities and processes associated with consciousness.

Reductive accounts that seek to *identify* experience with its observable physical correlates face the difficulty that such an identity would make experience a public object, contravening the privacy constraint. For example, Patricia Churchland, skeptical of any 'hard problem' of consciousness, suggests that qualitative experiences just *are* certain brain states (Blackmore, 2006, 60). In which case, since brain states are in principle completely physical and publicly accessible, so too must be experiences. But however closely you inspect, measure, and schematize the brain states associated with pain, you'll not find pain as an observable spatio-temporal property (Dennett, 1978). You'd have to instantiate (be) those states for pain to exist, and it would then exist only for you, which is not the case for your brain states. Jesse Prinz, a self-described physicalist, has developed an empirically grounded representationalist hypothesis that consciousness (at least in our case) is constituted by neurally-instantiated *attended intermediate representations* (AIRs) in the brain (Prinz, 2012). But in examining the neural AIRs (the vehicles) we wouldn't thereby see or access consciousness (the phenomenal content). The former are all in the public domain, the latter not. Michael Tye, another physicalist-representationalist, has proposed that consciousness consists of Poised, Abstract, Non-conceptual, and Intentional Content – PANIC (Tye, 2000). The underlying physicalist assumption of the PANIC hypothesis is that phenomenal content ultimately reduces to, and thus is identical to, some set of physically-instantiated representational goings-on – the vehicles. Or if a reduction isn't in the offing, then phenomenal content will have objective status as a categorically mental property of said vehicles. The first alternative places consciousness in the public domain, thus is ruled out on my view, while the second suffers from the problems confronting any sort of property dualism: explaining how consciousness as an objective mental property, something irreducible to the physical, emerges from a physically instantiated system and then (barring epiphenomenalism) goes on to play a causal role in behavior.

Susan Pockett proposes another type of phenomenal-physical identity, that consciousness is 'identical with certain spatio-temporal patterns in the electro-magnetic field' generated by the brain (Pockett, 2000). Were this the case, in observing and measuring those patterns, we would *per impossibile* be observing and measuring experiences as public objects. This applies to all hypotheses about consciousness which hold it to be objectively specifiable, even something as intangible (yet physical) as the collapse of the wave function in micro-tubules (Ekert et al., 1998). According to John Searle's biological naturalism,

> Everything that has a real existence has it in a single space/time continuum and the real existence of consciousness is in human and animal brains. Thoughts about your grandmother, for example, are caused by neuron firings and they exist in the brain as a feature of the system at a higher level than that of individual neurons. (Schneider and Velmans, 2017, 327).

Observing that higher-level feature, for instance some sort of neurally-instantiated functional organization, would be to observe the experienced thought, but this is ruled out by the non-objectivity of consciousness. Although Searle elsewhere acknowledges the subjectivity of experience – its 'first person ontology' – his insistence that only spatio-temporal entities can be real doesn't sit well with subjectivity.

6.2 Radical objectivisms

Panpsychism is among the more extreme approaches to reifying consciousness: the phenomenal is hypothesized to be a completely mind-independent property of the basic constituents of matter, such that experience or perhaps *proto*phenomenology (something quasi-phenomenal that can combine to constitute the phenomenal) might be present just about everywhere in spacetime (Strawson, 2006; Skrbina, 2009a). Thus far there is no empirical support for such proposals, and as noted above the evidence thus far strongly suggests that consciousness is associated with composite mind-systems doing particular representational, behavior-controlling and system-regulatory work. Panpsychism attempts an end run around the question of how qualitative content arises for mind-systems by positing the existence of *objective, system-independent* quanta of qualitative or pre-qualitative subjectivity. But this maneuver is effectively blocked, at least for the time being, by the absence of any proposal for, or evidence of, the way these strictly hypothetical quanta combine to become contentful phenomenal experiences for subjects like us (the "combination problem").

A variant of panpsychism, Russellian monism (RM) holds that the structural and dynamical regularities described by physics at the micro-level aren't all there is to the world; in addition there exist categorical natures which ground fundamental physical regularities, natures identical to the qualitative states of experiences (Alter & Nagasawa, 2015). In having experiences, we are directly acquainted with the intrinsic, objective, mind-independent self-nature of the world. This neatly solves two problems, that of naturalizing consciousness and that of identifying the *concreta* which some think lie behind the merely structural and relational characterization of the world given by physics. But like panpsychism, RM lacks any empirical support. It also runs afoul of the representational relation in supposing that phenomenal experience constitutes direct, non-representational but nevertheless epistemic contact with the (unrepresented) reality of intrinsic self-natures. Such contact may not be in the cards for the project of objectification since the representational relation can only afford our world-model a structural correspondence with reality (Ladyman and Ross, 2007). Moreover, the idea that physics needs supplementing by knowable intrinsic natures seems a holdover from folk physicalism: that reality must be ultimately concrete, partaking of some kind of stuff or substrate. What's mind-independently real, at bottom (should there be one), may not be under any such obligation.

An even more radical approach to objectifying consciousness is Riccardo Manzotti's 'spread mind' hypothesis, which holds that conscious experiences don't represent physical objects, rather they are *identical to* those objects that appear in consciousness:

> …to perceive something does not entail concocting a representation of something, but rather perceiving something means that the something is literally part of the experience of the subject. Whatever the subject sees is identical with a process beginning in the environment and ending in her brain. In turn, the perceived object would be identical to that process. There is no separation between the physical world and the experience of the subject. (Manzotti, 2011, 66)

On the assumption that the physical world is mind-independent, it is ordinarily thought to be intersubjectively accessible: we generally suppose that differently placed observers can observe and have experiences involving the same mind-independent object. However, since we each have distinct experiences of an object which may differ depending on our perceptual capacities and location, and since according to Manzotti objects are identical to those different experiences, on

his account there is no single, intersubjectively available mind-independent object to which we all have access. The claimed identity of experience and the physical world thus puts the presumptive mind-independent status of reality in doubt. Further, if there is no difference between the physical world and a subject's experience, this means my experience, and yours, ends up being publicly accessible, contravening the privacy constraint. Lastly, for Manzotti experiences can never be informative or misinformative, since they simply *are* what's physically the case, even dreams, hallucinations and visual illusions.[5] All told, the spread mind hypothesis is a tough sell for physicalists holding an internalist, supervenience-based view of consciousness, and even more so for representationalists like myself who suppose there's good reason to distinguish between mind-dependent representational content and what that content refers to or tracks, which is often a mind-independent physical state of affairs.

6.3 Better bets

Given my thus far pessimistic (and for reasons of space, incomplete) assessment of hypotheses concerning consciousness, what approaches might better conform to the constraints I have proposed? Stronger candidates would be those that accept the *prima facie* informational content of experiences *and* acknowledge that such content is proprietary, not an observable. It will come as no surprise, then, that I find Integrated Information Theory (IIT) (Tononi & Koch, 2014) to be a promising proposal, since it meets both these criteria in its attempt to explain consciousness. IIT wears its representational commitment on its sleeve, since information is ordinarily *about* something, or can be construed to be.[6] Secondly, in IIT the information as qualitatively rendered in experience only exists *for the system*:

> In IIT, information is meant to capture the "differences that make a difference" *from the perspective of the system itself* – and is therefore both causal and intrinsic. These and other features distinguish this "intrinsic" notion of information from the "extrinsic", Shannon notion… (Oizumi et al, 2014, p. 6, emphasis added)

Keeping in mind that by 'from the perspective of the system itself' IIT does not mean that the system is in an observational relation to its internal information (which would threaten an epistemic regress), but rather in an existential, *constitutive* relation (that is, of being), we can see that it respects the privacy constraint on consciousness. The central identity thesis of IIT, that 'an experience is identical with the maximally irreducible conceptual structure…specified by the mechanisms of a complex in a [system's] state' (Oizumi et al., 2014) means that experience exists only for the instantiating system, not as an observable such as the system itself.

I can't essay a detailed evaluation of IIT here but will simply note that certain of its implications seem intuitively implausible, for example that photodiodes host experience and that perfectly static systems might as well. Such intuitions will have to give way should the theory pan out, but the evidence in hand indicates that experience arises only in conjunction with complex systems that can engage in system-maintaining behavior with respect to their environments. Oizumi et al. and Tononi and Koch (2014) also speak from time to time of consciousness being 'generated' by

[5] 'It is our direct individual experience that is unerringly true; being one with the external world, it cannot be wrong.' – from Manzotti and Parks, 2018, http://www.nybooks.com/daily/2017/11/26/the-pizza-thought-experiment/.

[6] Oizumi et al. (2014) say: 'While emphasizing the self-referential nature of concepts and meaning [that is, of the system's informational content], IIT naturally recognizes that in the end most concepts owe their origin to the presence of regularities in the environment, to which they ultimately must refer, albeit only indirectly'.

physical systems; but if experience is *identical* to informational content – the maximally irreducible conceptual structure of a system – then it isn't generated as a further effect of being that structure. Another concern is that ITT does not explain why integrated information should feel like anything for the instantiating system, that is, be qualitative; but the theory is yet young.

A rather different take on consciousness, that of James Tartaglia (like myself, a former physicalist), recognizes that experience isn't found in the world as described by science. He says that '…the objective world lacks any evidence for the existence of experience: it is this fact, after all, that generates the problem of consciousness in the first place' (Tartaglia, 2016, 93). But instead of opting for illusionism or eliminativism, Tartaglia holds that experience is real *and* that it's representational (86-9): we perceive objects 'in virtue of having experiences' (88) so subjectively we are in indirect contact with the world, what he calls indirect realism. Each of us embodies a conscious perspective that bequeaths us a world given in terms of experience. However, he goes on to argue that experience itself *transcends* that world; it is part of a wider, final context of existence, of a transcendent being that can't ever be grasped as it is in itself (117-19).

I find Tartaglia's 'transcendent hypothesis' appealing since it gets so much right about experience, in particular its non-objectivity and representational nature. It also recognizes that ultimate reality cannot be captured in its self-nature precisely because we're always in a representational, perspectival relation to it. But I think it goes wrong in holding that consciousness is mind-*independently* real, that it participates in a transcendent reality. Tartaglia says '…the transcendent reality of experience is not actually caused by an objective world' (166), and so puts aside the question that most nags at naturalists: why, as the evidence suggests, does consciousness only attend the physical brain, doing only particular cognitive things? Even if we renounce *causal* explanations of consciousness, as I have, the obdurate, objective fact remains (a fact that should be included in our collective reality-model) that as far as we know only certain sorts of up-and-running systems host conscious experience; this cries out for a naturalistic explanation. To hold that consciousness is transcendent – a fundamental mind-independent reality – conveniently avoids that problem, much in the way that panpsychism avoids it by holding phenomenality to be something particulate and fundamental. The transcendent hypothesis thus marginalizes what seems to me a pressing philo-scientific question, that of explaining consciousness as a phenomenon that on its face is neither transcendent nor fundamental, but rather mind-dependent.

Thomas Metzinger takes on this question in his ambitious 'self-model theory of subjectivity' (Metzinger, 2003, 2009), elements of which were referred to in Section 2 above. The theory is explicitly representationalist, informed by and predictive of empirical findings, and seeks to explain the central conditions and features of phenomenal consciousness.[7] To be conscious is to instantiate a neurally-realized, functionally adaptive, self-in-the-world reality-model or representation, parts of which the system can't recognize as a model. On Metzinger's account, to take something in experience as real is to not grasp the content of experience *as a representation*: experience is thus transparent for the system (an invisible interface) and its content taken as reality (2003, 163-169). We consequently end up as naïve realists about at least some portion of the represented world, in particular ourselves *as concrete selves* (hence the 'self-model' theory).

[7] On the essential features of consciousness, see Metzinger (2003) pp.107-208.

Metzinger argues that the representational nature of consciousness is suppressed for biological systems such as ourselves as a matter of cognitive efficiency. To minimize energy consumption and afford real-time behavioral control, our representational capacities must be limited in their recursive (meta-representational) application, and this limitation can perhaps help explain why we end up hosting qualitative content. The adaptive (and perhaps logically necessary) *closing off* of what would be a paralyzing representational, epistemic regress means that the system will instantiate on the sub-personal level representational content that it can't *further* represent; Metzinger calls this 'autoepistemic closure' (2003, 131). Such content, in particular that associated with sensory channels such as vision, olfaction, hearing, and internal proprioception, perforce becomes cognitively impenetrable, an irreducible basic element of representation itself, that which the system can't specify any further component of or take as a representable object (except conceptually of course, as we're now engaged in). But what *is* this if not the essential characteristic of individual, basic qualia: their not-further-characterizable, non-decomposable, hence ineffable phenomenal character, the fact that in consciousness we can't get behind, or into, or away from such things as sweetness, primary red, pain, or any other basic unanalyzable sensory quality? And, according to Metzinger, since we're not in a position to take most qualitative content as representations (in particular that of most waking experience), qualities in experience are perforce taken as *object* properties, properties of things construed as mind-independently real.

This possible explanation of proprietary qualitative content as an outcome of being a necessarily limited behavior-controlling representational system, of which the above is just the briefest sketch, is among the most promising features of Metzinger's detailed and theoretically rich approach. Although he holds that consciousness is a biological phenomenon (2009, 58), it's notable that he offers no *causal* explanation of phenomenology nor any pat phenomenal-physical identity claim. Instead, he proposes a logical and adaptive *entailment* from the limitations of physically-instantiated representational functions to the existence of phenomenal experience for, and only for, the instantiating system. This, it seems to me, at least gets us in the vicinity of qualia construed as the irreducible basic elements of a private representational reality.

7. Phenomenal-physical parallelism

That said, I differ with Metzinger in his supposition that qualia, and the integrated, dynamic phenomenal gestalt of waking consciousness, can play a role in scientific accounts of behavior and the evolution of the neural processes that support consciousness (Metzinger 2009, 54-62). In these accounts, science can only deal in observables or what accepted theories suggest might eventually be observed or detected (e.g., dark matter, Hawking radiation). Since qualia and the experiences constructed out of them are not, and I think never will be, observables, they can't be cited as causal factors alongside or in addition to what their neural correlates accomplish in physical and functional accounts of behavior control. We therefore have to give up on objective phenomenal causation since the phenomenal, although real, doesn't appear in the objective world as represented by science. In the evolutionary cognitive arms race, natural selection undoubtedly selected for the behavior-controlling cognitive functions and underlying neural wetware associated with being conscious, but not for consciousness itself (Rosenthal, 2008; Oakley and Halligan, 2017). However, since we live our subjective lives completely within the ego tunnel of conscious experience, we unsurprisingly take experience to be a causal operator: I eat chocolate *because it tastes good to me*, right? Well, like the red of Dennett's afterimage, objectively that taste is nowhere to be found, so science will always default to the neural correlates of the taste in

explaining my chocolate habit. But in everyday life I will as a conscious subject continue unperturbed to cite the taste – a phenomenal particular – as the cause. We thus have two sorts of valid, predictive explanations, one involving my experience, another involving its physical correlates.

Objective phenomenal causation would require that experience somehow causally supplements the physical on the objective causal playing field, what we might call objective *explanatory space*. But again, we don't find my taste of chocolate in that space, that which involves physical spacetime. What is more plausible, or at least better governed by the constraints suggested above, is a non-interactive phenomenal-physical parallelism that involves respective subjective and objective explanatory spaces. Such parallelism respects the reality of both qualia (the private representational reality of experience) and its correlates (the collective *represented* reality of neural processes) without supposing they causally interact, which obviates the problem of phenomenal causation. As suggested above, it could be that some sort of non-causal entailment from physically-instantiated representational goings-on to qualia underwrites the reliable parallelism.

Keeping the phenomenal and physical on separate but correlated tracks also relieves us of the burden of somehow reducing qualitative feels to physical facts. Qualities, experiential and non-objective, are the not-further-specifiable basic terms in which the folk-physical world appears in consciousness; they aren't objects in that world that can be reduced to facts about more basic objects. Rather, folk-physical facts and truths about the world (e.g., that the apple in Section 4 is red and tastes sweet) are couched *in terms* of such qualities. Not having to concern itself with phenomenal-physical reduction, identity, or causation, physical science can proceed apace in developing an ever more predictive quantitative and conceptual grasp of reality.

8. Conclusion

As much as I've pleaded the case for phenomenal realism and the non-objectivity of qualia, the widely held naturalistic presumption of physicalism makes the conjunction of these claims at least counterintuitive and perhaps irredeemably obscure. Physicalists take composition, causation, reduction, and emergence to be primary explanatory relations among phenomena, so of course want to apply them in naturalizing consciousness. Conscious mind-systems are natural, objective, physical phenomena, resident in spacetime, so if consciousness is real must it not also be discoverable there? The idea that each person's conscious experience constitutes a private representational reality will of course seem suspect under a regime that requires all of what's real to be objective. But if content, in particular phenomenal content, can't be objectified, then consciousness won't be discoverable in spacetime, and indeed that's the current state of play: experience is a private, not public affair, and explanations must somehow respect this (objective) fact about consciousness, at least until it's shown that there is no such fact, or, as illusionists hold, there is no such thing as phenomenal experience.

The way forward, I'd suggest, is to continue full steam ahead with the investigation of the neural and functional correlates of consciousness; this will inform a *science of representation* that may lend plausibility to the possible entailment from certain types of representational processes to experience. The nature of that entailment, should it exist, may become clear as Metzinger's self-model theory and its representationalist competitors and collaborators, e.g., IIT, Prinzian AIRs, and Tye's PANIC account, are refined in light of further research. Understanding the biological and artificial architectures of world-responsive informational systems, as for instance in the

predictive coding paradigm, may show how instantiating a suitably ramified and integrated reality-model makes representational content qualitative for the system – not as a matter of causation or emergence but of *representational necessity*. We may not find a satisfying *physicalist* account of how consciousness arises in nature, since it doesn't appear in the physical world, but we might find sufficient resources in representationalism to close the explanatory gap (Levine, 1983).

We can understand physicalism as a well-intentioned and natural attempt at cognitive unification – a global metaphysical thesis about the fundamentally physical nature of reality – but consciousness puts physicalism under considerable pressure.[8] Physicalism as a metaphysical thesis forgets, perhaps, that the world as objectified in folk-physicalism and science is a *represented* reality, a world-model, not unrepresented reality. A better, more self-consistent global naturalistic realism will incorporate the indispensable role of representation in the collective world-model itself. Since our world-model represents what we take to be real – our ontology – as thus expanded it will assert the reality of representation, not as a separate substance or property, but as a condition of being knowers. If the representational relation puts us as knowers at one remove from unrepresented, mind-independent reality, as natural, physical creatures we ourselves are nevertheless situated in that reality. This means our conscious experience too, even if not discoverable in spacetime, is fully within the natural world as it engages, through us, in the project of self-knowledge.

Acknowledgements

My thanks to anonymous referees for their helpful comments and to those who provided feedback on earlier drafts or sections of this paper, including Susan Blackmore, Peter Clark, Daniel Dennett, David Duffy, Peter Hankins, Nicholas Humphrey, James Tartaglia, and Paul Torek.

References

Alter, T. & Nagasawa, Y. (2012) What is Russellian monism?, *Journal of Consciousness Studies,* **19** (9), pp. 67-95.

Baars, Bernard J. (1997) Contrastive phenomenology: a thoroughly empirical approach to consciousness, in Block, N., Flanagan, O. & Guzeldere, G. (eds.) *The Nature of Consciousness: Philosophical Debates,* Cambridge, MA: MIT Press.

Blackmore, S. (2006) *Conversations on Consciousness*, Oxford: Oxford University Press.

Bourget, D. & Chalmers, D. (2018) eds. *PhilPapers*. https://philpapers.org/browse/theories-of-consciousness.

Clark, A. (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science, *Behavioral and Brain Sciences,* **36** (3), pp. 1-24.

[8] For acknowledgement by physicalists of the limitations of physicalist explanations of consciousness, see Kim, 2005 and Howell, 2013. Kim accepts that qualia can't be functionalized, thus they resist any straightforward incorporation into physical science, while Howell argues that phenomenal consciousness, although ultimately physical, can't be objectified, hence we should be 'subjective physicalists'.

Clark, T. (2005) Killing the observer, *Journal of Consciousness Studies*, **12** (4-5), pp. 38—59.

Dehaene, S. & Changeux, J. P. (2011) Experimental and theoretical approaches to conscious processing, *Neuron* **70** (2), pp. 200-227.

Dennett, D. C. (1978) Why you can't make a computer that feels pain, *Synthese*, **38** (3), pp. 415-456.

Dennett, D. C. (1987) *The Intentional Stance*, Cambridge, MA: M.I.T. Press.

Dennett, D. C. (1998) The myth of double transduction, in Hameroff, S., Kaszniak, A., & Scott, A. (eds.) *Toward a Science of Consciousness II, The Second Tucson Discussions and Debates*, Cambridge, MA: MIT Press.

Dennett, D. C. (2015) Why and how does consciousness seem the way it seems?, in Metzinger, T. & Windt, J. M. (eds.) *Open MIND*: 10 (T). Frankfurt am Main: MIND Group.

Dennett, D. C. (2016) Illusionism as the obvious default theory of consciousness, *Journal of Consciousness Studies*, **23** (11-12), pp. 65-72.

Dennett, D. C. (2017) *From Bacteria to Bach and Back: The Evolution of Minds*, New York: W.W. Norton & Company.

Ekert, A., Jozsa, R., Penrose, R., & Hameroff, S. (1998) Quantum computation in brain microtubules? The Penrose–Hameroff 'Orch OR' model of consciousness, *Phil. Trans. R. Soc. Lond*, 356, pp. 1869–1896.

Feinberg, T.E. & Mallatt, J. M. (2018) *Consciousness Demystified*, Cambridge, MA: MIT Press.

Frankish, K. (2016) Illusionism as a theory of consciousness, *Journal of Consciousness Studies*, **23** (11-12), pp. 11-39.

Gamez, D. (2014) The measurement of consciousness: a framework for the scientific study of consciousness, *Frontiers in Psychology*, **5**. 714.

Ginsburg, S. & Jablonka, E. (2019) *The Evolution of the Sensitive Soul*, Cambridge, MA: MIT Press.

Goff, P. (2016) Is realism about consciousness compatible with a scientifically respectable world view?, *Journal of Consciousness Studies*, **23** (11-12), pp. 83-97.

Howell, R. J. (2013) *Consciousness and the Limits of Objectivity*, Oxford: Oxford University Press.

Kanwisher, N. (2001) Neural events and perceptual awareness, *Cognition*, **79** (1-2), pp. 89-113.

Kim, J. (2005) *Physicalism, or Something Near Enough*, Princeton, NJ: Princeton University Press.

Ladyman, J., & Ross, D. (2007) *Every Thing Must Go: Metaphysics Naturalized*, Oxford: Oxford University Press.

Levine, J. (1983) Materialism and qualia: the explanatory gap, *Pacific Philosophical Quarterly*, **64**, pp. 354–361.

Lewis, C.I. (1929, 1991) *Mind and the World Order: Outline of a Theory of Knowledge*, New York: Dover Publications.

Manzotti, R. (2011) The spread mind: is consciousness situated?, *Teorema*, **30** (2), 55-78.

Manzotti, R. & Parks, T (2018) On Consciousness, *New York Review of Books*, http://www.nybooks.com/topics/on-consciousness/

Metzinger, T. (2003) *Being No One: The Self-Model Theory of Subjectivity*, Cambridge, MA: M.I.T. Press.

Metzinger, T. (2009) *The Ego Tunnel: The Science of Mind and the Myth of the Self*, New York: Basic Books.

Oakley, D., & Halligan, P. (2017) Chasing the rainbow: the non-conscious nature of being, *Frontiers in Psycholo*gy, **8** 1924.

Oizumi, M., Albantakis, L., & Tononi, G. (2014) From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0., *PLOS Computational Biology* **10** (5).

Pockett, S. (2000) *The Nature of Consciousness: A Hypothesis*, Lincoln, NE: Writers Club Press.

Prinz, J. J. (2012) *The Conscious Brain: How Attention Engenders Experience*, Oxford: Oxford University Press.

Revonsuo, A. (2015) Hard to see the problem?, *Journal of Consciousness Studies*, **22** (3-4), pp. 52-67.

Robinson, W. S. (2010) Epiphenomenalism, *Wiley Interdisciplinary Reviews: Cognitive Science* **1** (4), pp. 539–547.

Rosenthal, D. M. (2008) Consciousness and its function, *Neuropsychologia*, **46** (3), pp. 829-40.

Schneider, S. & Velmans, M. (eds.) (2017) *The Blackwell Companion to Consciousness*, Hoboken: Wiley and Sons, Ltd.

Sellars, W. (1962) Philosophy and the scientific image of man, in Sellars, W., *Science, Perception and Reality*, London: Routledge & Kegan Paul.

Seth, A., Suzuki, K., & Critchly, H.D. (2011) An interoceptive predictive coding model of conscious presence, *Frontiers in Psychology*, **2** 395.

Skrbina, D. (ed.) (2009a) *Mind that Abides: Panpsychism in the New Millennium*. Philadelphia: John Benjamins.

Sterling, P. & Laughlin, S. (2015) *Principles of Neural Design*, Cambridge, MA: M.I.T. Press.

Strawson, G. (2006) Realistic monism: why physicalism entails panpsychism, *Journal of Consciousness Studies*, **13** (10-11), pp. 3-31.

Tartaglia, J. (2016) *Philosophy in a Meaningless Life: A System of Nihilism, Consciousness and Reality*, London: Bloomsbury.

Tononi, G. & Koch, C. (2014) Consciousness: here, there but not everywhere, *arXiv.org*, arXiv:1405.7089.

Tye, M. (2000) *Consciousness, Color and Content*, Cambridge, MA: MIT Press.

Van Gulick, R. (2014) Consciousness, in Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy*, Spring 2014 edition, pp. 34-43.