



Our “Cognitive Limitations” and the Hard Problem of Consciousness

ABSTRACT: *Philosophers have conjectured that human cognitive limitations might preclude our ever resolving the hard problem of consciousness. Few, however, have offered suggestions as to what it might be about our conceptual apparatus that poses the problem. I do so in this essay, arguing that our central difficulties lie with two conceptual categories that pervade philosophical discussion of the hard problem. They are compositional concepts – part/whole, constituents/constitution, and the like – and instantiational concepts: properties/objects, universals/particulars, and the like. I look at the uses of these two conceptual categories in four contexts in which the hard problem is considered: multiple realizability, zombies, mental causation, and panpsychism. I show that the two conceptual categories run into the same kinds of obstacles in each case, which suggests that they might be key to the cognitive limitations about which some have speculated*

KEYWORDS: hard problem of consciousness, multiple realizability, zombies, mental causation, panpsychism

Four debates—over multiple realizability, zombies, mental causation, and panpsychism—have assumed prominence in recent decades in the philosophical attempt to grapple with the hard problem of consciousness. There are others, of course, but these four are prominent perhaps in part because they form a kind of rough symmetry. Multiple realizability invites us to consider how any given mental state might possess more than one underlying physical substrate. Zombies beckon us to ponder how any given physical substrate might underlie more than one mental state—or, more exactly, a mental state on the one hand and the lack of a mental state on the other. Mental causation requires us to comprehend how the mind might descend into the physical plane. Panpsychism beckons us to comprehend how the physical might ascend into the mental plane. Together, the four debates stake out much of the central terrain in our effort to explain how the mind can so thoroughly depend on the brain, and yet seemingly remain so utterly distinct from it: the hard problem. None, though, has reached any kind of consensus.

These continuing controversies reaffirm the conjecture that, borrowing from Thomas Nagel (2012:128), the answer to the hard problem might lie “beyond our reach, in virtue of our intrinsic cognitive limitations” and “the inadequacy of our present concepts.” What concepts? Although Nagel’s view is a popular one, there has been little exploration of what exactly it is about our cognitive concepts that

could place the hard problem of consciousness beyond our grasp. Taken together, though, the four debates provide rich material for beginning to answer *that* question.

Across the four debates, two classes of conceptual apparatus do much of the heavy lifting. One class is compositional: part/whole, constituents/constitution, and the like (I here follow Gillett 2016:33; Audi 2012:667, and others in treating constitution, at the broad level I am exploring here, as a compositional relationship). The other class is instantiational: properties and objects, universals and particulars, and the like. Sometimes what appears in one class might appear in the other: properties can be instantiated and they can compose other properties. But the two relationships remain distinguishable.

The differences between these two broad classes of conceptual apparatus, and the shadings within them, have been subjected to considerable analysis. I abstract away from much of that. My interest is more foundational. I examine the usage of these two basic conceptual classes, compositional and instantiational, in several attempts to advance the debates over multiple realization, zombies, mental causation, and panpsychism, and to explore whether, and how, they might create the conceptual limitations to which Nagel refers. More hopefully, even though we are far from resolving the mystery of consciousness in a way that commands consensus, I ask whether our usages of these two conceptual relationships point to what might be required for us to take at least one important step toward such a resolution. Can we gain at least some clarity on what it is we're looking for – the kinds of limitations we have to break through – as we seek to resolve the hard problem?

McGinn (1989) has argued, of course, that the hard problem might lie beyond our cognitive reach because our introspective and perceptual capacities, together, do not suffice to reveal the mind-body connection. I take no issue with McGinn's argument, since we may well face more than one kind of limitation on our capacities to resolve the hard problem, although McGinn focuses comparatively more on perceptual than conceptual issues. In any event, I am simply advancing another candidate, having to do not with introspection and perception but our concepts of composition and instantiation.

1. Composition and Instantiation

At the core of those limitations, I argue, is a kind of mirror imagery that composition and instantiation display. Each supplies something the other lacks. Think of entities that relate compositionally, such as parts and wholes. The compositional relationship readily accommodates their differences in qualities. The parts of a car—the engine, tires, chassis, etc. —and the whole, the car, that they make up can have very different qualities: a tire is rubbery and round, a car is sleek and comfortable. Likewise with a lump of clay and the statue it constitutes. As Gillett (2007a:84) says, “composition usually relates qualitatively different kinds of entity.”

But while the parts of a car and the whole they compose both belong to the same ontological category—namely, objects—the compositional relationship faces difficulties when parts and whole belong to the different ontological categories of object and property: when the parts fall into the category of objects—such as the

brain or its neurons—and the whole falls into the category of properties, such as consciousness or a mental state (Clapp 2001; Antony 1999). Compositional relationships do not readily bridge the ontological difference between objects (as parts) and properties (as wholes); generally, if the parts are objects, then the wholes that they compose must be objects too. On some ontologies, perhaps, this need not be the case. In the four debates, however, any attempt to suggest otherwise would become problematic, replacing the hard problem of consciousness with the hard problem of how objects can compose properties.

Certainly, it’s often said that it’s the brain’s physical *properties*—such as electric charge—and not simply those objects, the brain itself or its neurons, that compose mental properties such as pain. But even when the physical and mental are both described as properties, composition—as I shall discuss—still faces an ontological divide. Compositional relationships cannot, for many critics, convincingly account for the ontological distinctiveness—the ontological reality—of mental properties from physical ones, in a way that satisfactorily addresses the hard problem.

The instantiational relationship, meanwhile, can of course account for the difference in ontology between object (brain) and property (mental states). Any particular object on the one hand, and the properties—or universals—it instantiates unproblematically differ in their ontological categories. Engines are objects, while engineness is a property. But the instantiational relationship itself cannot account for a difference in qualities between the instantiating object and the instantiated property, in this sense: Properties must be “wholly present in the things” —the objects—“that accommodate them” (Kroedel 2020:24). Any particular engine, the object, must resemble the property of engineness that it instantiates, such that we can identify the property in the object. In other words, both object and property in some sense share an engine quality. And yet the “soggy” and “grey” qualities of that object, the brain, and the “technicolor” qualities of properties such as mental states (McGinn 1989:349), differ profoundly. While a compositional relationship can in principle accommodate differences in qualities, an instantiational one seemingly cannot.

Certainly, it’s often said that it’s the brain’s physical *properties*—such as electric charge—and not simply those objects, the brain itself or its neurons, that instantiate mental properties such as pain. But again, the qualitative difference between the two types of properties—the instantiating physical and the instantiated mental—remains, making it difficult for instantiation to account for their relationship. I note here that in what follows I will, for ease of exposition, speak simply of physical objects themselves such as the brain or its neurons, and not their physical properties, like electrical charge, as instantiating mental properties, like pain. Nothing is lost in doing so for purposes of my discussion. As Macdonald (2005:258) says, “an instance of a universal just is the thing that has it, that is, a particular concrete object. Thus, for example, an instance of the property, red, just is the red bird.” I will follow that usage here.

A couple of other terminological points: First, I use the term “quality” to capture what it is—within a typical instantiational relationship—that a property like “engineness” shares with the particular objects, the engines, that instantiate it. They share an “engine” quality. They share, in other words, something that

allows us to say that such and such a property is instantiated in such and such an object; one can recognize the qualities of the instantiated property in the object that instantiates it. Perhaps we might want to say that an engine itself has the quality of being an engine, and the property of engineness has the quality of making something an engine, but both of those qualities in turn rely on the idea of “engine,” and that is what I am referring to as their common quality here.

Some might say that engineness and engines share an enginey property instead of an enginey quality. But that would require the awkwardness of our talking about properties of properties—of saying that the property of engineness and the particular objects known as engines share an enginey property. Here, following Aizawa and Gillett (2009:188), I will simply use “qualities” when I am talking about what the two partners in an instantiational relationship must share—or what the various levels in a compositional relationship needn’t share. What will matter, when it comes to qualities, is the way in which instantiational conceptualization requires instantiating objects to share qualities with their instantiated properties, just as what will matter for compositional conceptualization is the ease with which it allows for composing and composed entities to display different qualities.

Second, I use the term “instantiation” instead of “realization.” Many philosophers treat the two terms synonymously. But many others speak of composition too, and not just instantiation, as a “realization” relationship: the parts of an entity realize the whole, just as an object realizes a property (Bennett 2017:9). I will thus use the term “instantiation,” not realization, to distinguish it from composition. As Oliver (1996:15) says, instead of “realizes,” philosophers “like to use the fancy word ‘instantiate.’”

In what follows, I do not take any position in the four debates under discussion, nor on the various claims I discuss. I do not make any argument *in* the debates over multiple realizability, zombies, mental causation, and panpsychism; rather I make an argument *about* the debates themselves. I reframe them through a critical lens: one that focuses on how the two conceptual distinctions, compositional and instantiational, work in each of the four. It is noteworthy that the four debates, though in many ways very different, all nevertheless share this same central feature: The presence of both compositional and instantiational argumentation, and the ways they conflict and relate, are remarkably similar in each. All of which suggests that perhaps herein lies a possible clue to at least some of our conceptual issues as we wrestle with the hard problem.

I want to emphasize that—given that I am looking at some basic commonalities across four wide-ranging debates—my uses of those terms, compositional and instantiational, are necessarily broad-gauge. They are meant to call attention to some basic features that the four debates share. In all of these debates, when the relationship between brain and mental states is elaborated compositionally, any difference in qualities between them is more readily explicable; any difference in ontology is not. Likewise, when the relationship between brain and mental states is framed instantiationally, any difference in ontology between them is readily explicable; any difference in qualities still invites explanation. It is the attempt to wrangle these two basic features of composition and instantiation, within the four debates, that is the focus of my discussion. Other ways of distinguishing between

various kinds of composition, or various kinds of instantiation, which might be appropriate for other purposes, lie beyond the topic I address.

I will claim that the two conceptual relationships display much the same broad pattern across the four central debates as they are used to wrestle with the hard problem. In fact, it sometimes seems as if composition is there to catch instantiation when it falls, as instantiation is there for composition when it falters. In the ways in which each makes up for the other’s shortcomings, they indicate more precisely what the nature of our cognitive limitations might be.

I emphasize that I am not talking about all of the debates the hard problem has instigated, just the four central ones of multiple realization, zombies, mental causation, and panpsychism. Nor am I making a statement about all of the many arguments that philosophers have mounted even within those debates, just about one way of understanding some foundational conceptual limitations that many confront. My discussion, too, is meant to be suggestive, not exhaustive, since I can here look at only some of the ways in which compositional and instantiational concepts are used in the four debates. Nor, finally, am I claiming that the “hard problem” would be resolved if we simply overcame the conceptual limitations posed by composition and instantiation; much more will surely need to be done. But given the commonality of the issues they face in the four debates, overcoming them in some way might well be a necessary step.

2. Multiple Realizability, Zombies, and Compositional Conceptualization

Multiple realizability—the claim that a single given mental state can be multiply realized in any number of different physical substrates—has readily lent itself to compositional conceptualization. Take the mental state of headache pain, with its qualities of pounding, smarting, and throbbing, and its multiple possible physical substrates, such as collections of carbon or silicon or metal parts, each of them displaying their own, very different qualities such as lumpishness, luster, or magnetism. In an analogous way, any given composition—any whole, any constitution, and the like—can also easily display qualities that differ substantially from those of its various possible collections of parts or constituents. A car—with certain qualities of sleekness, agility, and speed—can be composed of numerous diverse collections of parts, each of them displaying their own, very different qualities of power, flexibility, and vulcanization, depending on whether the collection is a battery engine, disc brakes, and convertible roof, or a combustion engine, drum brakes, and metal roof, and so forth.

Of course, much depends on how we understand “pain” or, for that matter, “car.” But my concern here lies in explicating why compositional conceptualization has proved ready-to-hand for the philosophical effort to account for multiple realizability. If “wholes often,” as Michael Silberstein (1998:480) says, “possess” qualities “that their parts do not” —and even as wholes are entirely dependent on those parts—then a compositional relationship can explain at least how mental states could possess qualities that differ from those of their physical substrates, even as mental states are entirely dependent on

their substrates. Specifically, it accounts for how a mental state can possess the same qualities when its various physical substrates display a variety of qualities.

It is such a difference in qualities that allows Aizawa and Gillett (2009), for example, to draw on compositional analysis to explicate the conceptual possibility of multiple realizability. Describing the relationship between various physical substrates and a mental state as one of “parthood” or “constitution,” Gillett (2007a:84) uses the compositional relationship between carbon atoms and a diamond as an analogy:

composition usually relates qualitatively different kinds of entity. For example, diamonds have [qualities] like hardness, and hardness contributes the powers to scratch glass. But no carbon atom has the [quality] of hardness, nor any property that contributes the power to scratch glass.

If mental states are composed by physical substrates – or as Gillett (2007b:196) says, “psychological states. . .are composed by brain areas”—in the same way that diamonds are composed by carbon atoms, then we are able to grasp how the mental, and the physical that composes it, can be so “qualitatively different.” Since multiple realizability denies one-one identity between mental states and physical substrates, it is sustained by any framework that allows us to distinguish the qualities of a mental state—its sadness or vividness—from those of its various physical substrates, whether neurons, silicon, or metal, as a compositional relationship does. The qualities of the composed—the mental state—can differ from those of the composing physical substrates, such that the qualities of the mental state can remain the same while those of the substrates can vary. As Polger and Shapiro (2016:108) put it in using compositional conceptualization to characterize Aizawa and Gillett’s position, a mental state is “multiply realized because there are variations in its parts.”

Compositional conceptualization assumes a role as well in the debate over zombies. Consider an argument Eric Olson (2018) makes that supports the conceivability of zombies. Think of a composition such as Henry Moore’s statue *Reclining Figure*. It possesses qualities that differ from those of the lump that composes it, namely bronze ingots fused together, even though the constituted sculpture and the constituent lump are physically identical. We attribute qualities such as languorousness, impassiveness, and repose to the statue—the constituted entity—and qualities such as hardness, matte finish, and undulation to the lump of bronze, the constituent entity. The lump can lose some of its qualities without the statue losing its qualities, and vice versa (see generally Rea 1995). If we attributed languorousness, impassiveness, and repose to the lump as well as to the statue it composes, Olson suggests, then we would have not one but two entities before us that possess those qualities—the statue and the lump—and that would be nonsensical.

Likewise, Olson argues, a body and the conscious person it physically constitutes—or, for purposes here, a brain and the conscious mind it physically constitutes—have different qualities. Your brain possesses chemical and electrical

qualities, for example, that differ utterly from your conscious mind’s qualities of pain or passion. As the compositional analogy with the lump and the statue suggests, even though the brain and the conscious mind are physically identical, we would not say that there are two entities here each possessing the qualities of pain or passion, your brain and your mind. And so we have allowed that your very own brain can possess physical qualities such as chemical snapping and electrical sparking without the mental qualities, the pain and the passion, possessed by the mind it composes. Your very own brain, then, Olson says, is a zombie. And so we have allowed for the conceptual possibility of zombies.

Olson draws various conclusions about the debate over zombies from his analogy (see also Kripke 1980:145) of the brain-mind relationship to the constituent-constituted lump-sculpture relationship. One of them is precisely that if brain relates to mind as constituent does to constituted, then a zombie brain is conceivable. Such compositional conceptualization, for Olson, allows for differences between the qualities of constituents and what it is they constitute, in ways that make the brain conceivable without the conscious mind it composes. Whether conceivability then implies possibility is, of course, another voluminously debated matter (see, e.g., Aydede and Güzeldere 2005). But since my concern here is with whether our basic conceptual apparatus might bar our ever bridging the explanatory gap, conceivability is the question at issue. It is with the use of compositional conceptualization that we first get what Chalmers calls a “positive” (imaginative) as opposed to a merely “negative” (logical) apprehensive purchase on how, even though consciousness might depend utterly on the brain, the brain need not guarantee the existence of consciousness: conscious mental states are something over and above it.

2.1 Compositional Conceptualization’s Difficulties with Multiple Realizability

In establishing beachheads for multiple realizability and zombies, compositional conceptualization has proved useful. A question, though, arises. Parts and the wholes they compose—or constituents and the constituted entity they compose—can certainly differ in their qualities. But with certain exceptions not relevant here, the composing and composed entities should share the same ontology such that if the former are objects, then the latter should be objects as well, not properties.

Certainly, in the analogies used by Gillett, Olson, and others, composed and composing do share the same object ontology. Carbon atoms—parts—and the diamond, the whole, might differ in their qualities, but all of them, parts and whole alike, are ontological objects. The same, too, with a statue and the lump that composes it. However different in their qualities, they each—the composed and the components—are objects. Mental states and their physical substrates, however, possess different ontologies; specifically, mental states are properties while neurons or silicon chips are objects. This poses a difficulty for compositional conceptualization as it’s sometimes used in both the multiple realizability and the zombie debates.

Consider multiple realizability first. It’s hard to see how objects like “brain areas” could compose a property such as a “psychological state” (Gillett 2007b:196).

That's so even if compositional conceptualization enables us to see how composed and composing can profoundly differ qualitatively, such that the same mental state can be conceived to have many different physical substrates. Moser and Trout (1995:210) make this point in discussing multiple realizability, noting that while “[s]patiotemporal concreta can be composed, or made up, of physical constituents, . . . multiply-realizable universals . . . are not properly regarded as composed of physical constituents.” Composition, then, runs into difficulty when the composing entities are “physical constituents”—ontological objects such as brain areas—and the composed entities are ontological “universals” (or properties) such as psychological states.

Of course Aizawa and Gillett (2009:187, 199)—along with others—speak not just of objects like “areas in the brain” but of the brain’s “biochemical properties” as composing a “psychological property” like the mental state of pain. Here, the composing and composed entities do share the same ontology; both are properties. And yet there still remains an ontological difference for composition to bridge. For the question remains as to whether composition can accommodate the necessary “ontological distinctiveness,” as Wilson (2021:20) terms it, between composed mental properties and component physical properties: a distinctiveness necessary, in other words, to bar a reduction of the mental to the physical that would undermine multiple realizability.

As Polger and Shapiro (2016:39) note, for example, “the sort of variation in composition that interests Gillett and Aizawa is, by their own account, no obstacle to reduction; so it is not suited to do the work that . . . multiple realization is supposed to do.” Composition of mental by physical properties might account for the qualitative differences—but it cannot account for the ontological distinctiveness—between the composed mental and the various component physical properties that makes multiple realizability conceivable. Gillett (2007b:203, 206) himself notes this. The “Argument from Composition,” he observes, is thoroughly “ontological in nature, reduc[ing]” composed entities like mental states “down to an ontology encompassing” component physical entities, and “entail[ing] that we should only accept the existence of component entities.” An “ontological reductionist,” Gillett (2007b:204) continues, can “argue that we should not be committed to both component and composed entities.”

If compositional conceptualization puts up “no obstacle” to ontologically reducing composed mental properties to component physical properties, then—for Polger and Shapiro—it cannot account for a sufficient ontological distinctiveness between composed and components to allow for multiple realizability. (Gillett himself advances a different, “machretic” compositional model but remains uncertain as to whether it applies to the debate over consciousness; see Gillett 2016:42). More generally, it's because composition faces difficulties lending a distinct ontological reality to composed mental properties that it needs to be supplemented with unspecified metaphysical or nomological principles or laws (see Chalmers 2008). Such supplementary laws or principles are necessary to explain how physical properties could compose mental properties that have a sufficiently distinct ontological reality so as to bar their reduction. But then it's those laws and principles—and not composition itself—that would do the heavy

conceptual lifting required to bridge the ontological distinctiveness between component physical properties and composed mental properties.

Absent such supplemental laws or principles, when composition alone is said to relate physical properties (as parts) to mental properties (as wholes), it runs into critics who will point out that “what is mereologically higher is not ontologically higher” (Rudder Baker 2019:210). Or, as Chalmers (2008:251) notes, the mere differences in qualities that composition allows for between composing and composed is not, by itself, sufficient to accord a distinct ontological reality to the composed, or to bar its ontological reduction to its components. Whether the components are deemed to be physical objects or physical properties, then, they cannot—for critics—accommodate the ontological differences they must bridge if they are to compose mental properties in a way that accounts for multiple realizability.

2.2. Compositional Conceptualization’s Difficulties with Zombies

Turn now to the ontological difficulties composition faces in the case of zombies. Certainly, Olson’s use of compositional (constitutional) conceptualization does account for how the brain and the conscious mind could possess very different qualities, as do lump and statue, making our very own brain conceivable as a zombie. And it might also seem as if, insofar as the brain can compose (as Olson says) a “conscious person” or a conscious mind—in other words, a person or mind that instantiates consciousness—we would avoid the ontological difficulty, since both brain and the mind (or person) are objects: the one physical, the other nonphysical.

But we then run into the matter of how to conceive of that nonphysical object, the mind (or person), as instantiating the property of consciousness or, more specifically, the mental properties of pain and so forth. How can we locate the qualities of a property like pain in the mind, a nonphysical object, that we can’t even apprehend? We can’t identify the qualities of that property, pain, in that nonphysical object in order to establish that instantiation, since we can’t identify that nonphysical object itself in the first place. That’s no less difficult than trying to locate the qualities of a property like pain in that physical object, the brain itself.

We thus return to the question of how the brain, an object, can itself directly compose consciousness or mental states, which are properties, and not simply compose other objects—minds—which instantiate those properties. Indeed Olson at times refers to the entity the brain composes in “property” terms—he calls it “consciousness” or “being conscious” —instead of in object terms like a “conscious mind” or a “conscious person”. Compositional (constitutional) framing cannot, though, bridge the ontological gap between an object like the brain, and properties like consciousness or mental states.

Olson himself goes on to note that a compositional approach still invites the question of how a property like consciousness could be constituted by an object like the brain, given the ontological difference between them. We still need, Olson (2018:223) says, to “bridge. . .the gap between a thing’s physical nature and its [also] being conscious.” We need, Olson continues, a further “metaphysical

principle linking consciousness with the conjunction of [that] physical nature *and* the right metaphysical kind” (italics mine) —or what I would here call “ontological kind.” That further principle must, Olson elaborates, be “something of the form: Necessarily, a thing is conscious if and only if it has physical nature N *and* belongs to metaphysical kind K” (italics mine).

In other words, a thing is conscious only if it’s constituted by a certain physical “thing” or object, the brain, and over and above that, somehow belongs to a different ontological kind than the physical entity that constitutes it. Compositional (constitutional) conceptualization by itself, as Olson acknowledges, cannot bridge—without a further unspecified “principle” —the ontological difference between mental and physical that partially motivates the hard problem; it can only explicate the qualitative differences between them that make zombies conceivable. Compositional conceptualization itself cannot explain how mental states come to accompany a brain, given the ontological differences between them. But of course compositional conceptualization does make it possible—because of the qualitative differences it accommodates between mental states and brain—for us to conceive of the same brain as both accompanied and unaccompanied by mental states.

3. Multiple Realizability, Zombies, and Instantiational Conceptualization

It is perhaps because compositional conceptualization runs into these difficulties that instantiational conceptualization can be found just as often in attempts to flesh out both multiple realizability and the possibility of zombies. Instantiational concepts can readily account for how mental and physical could possess the divergent ontologies of property and object, even as the one totally depends on the other. Consider, as Polger and Shapiro invite us to do in discussing multiple realizability, the idea of “corkscrewness,” which they describe as a “property” that requires “particular objects” —the innumerable corkscrews in the world such as A’s cantilevered one, B’s two-pronged one, or C’s old-fashioned one—to instantiate it.

By analogy, the mental state of pain is a property that requires its various physical substrates in the material world, objects such as my C-fibers, or a lizard’s clump of neurons, or a robot’s chips, to instantiate it. Instantiation then could—and for many it does—provide the basic tools for conceptualizing multiple realizability. And if a mental state, like a headache, is a property—or a universal (Fisher 2018:218; Williams 2007:199) —then that would explain how it can be instantiated in any number of different objects, of physical substrates (Polger and Shapiro 2016:28; see also Moser and Trout 1995:212; Macdonald and Macdonald 1986:149). Instantiation begins to explain how a property like a mental state could be something distinct from any particular object, whether a brain or something else, that instantiates it, just as corkscrewness is from any particular corkscrew that instantiates it. In applying instantiational conceptualization, Polger and Shapiro of course remain skeptical about many cases that others have offered to exemplify multiple realizability. What they do is explicate what an instantiational take on multiple realizability must look like.

A problem, though, remains. Properties and the objects they depend on to instantiate them do certainly differ in their ontologies. Properties and the objects that instantiate them are, however, supposed to share at least some similar qualities, as the corkscrewishness/corkscrew analogy suggests. The property of corkscrewishness and my friend’s particular object, their particular corkscrew, that instantiates it at least share, in their different ways, a corkscrew-y quality. But instantiatonal conceptualization cannot capture the qualitative difference between brain and mental states—even though it can capture the ontological difference—because instantiation relies on at least some qualitative similarity between instantiated and instantiator of a sort that doesn’t exist in the brain/mental state case.

As countless contributors to the debate over the hard problem have pointed out, you can look as hard as you want in the brain. You won’t identify in that physical object, with its qualities of substantiality and sponginess, any of the qualities, any of the dolor or discomfort, of the property, the mental state of pain, that it is supposed to instantiate (Antony 1999:22; Doggett and Stoljar 2010:132). For an object to instantiate a property, it must display—in some fashion—something, which I am calling (with Gillett) a “quality,” identified with that property. And so instantiation encounters difficulties in its efforts to conceptualize the mind-brain relationship and explain multiple realizability, the converse of those faced by compositional conceptualization, which can capture the qualitative differences between mind and brain but not the ontological differences.

To explore this further, consider the use of instantiation in the case of zombies; Chalmers (1996:146–50), for example, employs the concept of “instantiation” in his account of their conceivability. Normally, as I have suggested, instantiation requires some similarity in qualities between an object and the property it instantiates. The relationship between the brain and mental states would seem to uniquely depart from this requirement of instantiation.

As it turns out, though, zombie proponents rely precisely on that unique departure, even as they make an instantiatonal case for the conceivability of zombies. That’s because the instantiatonal relationship between mental states and the brain differs in a key way from every other instantiatonal relationship between a property and its instantiating objects, such that we should not—in the mental state/brain case—expect to identify the qualities of the property in the object.

Nagel (1974), for example, sees mental states as “properties” instantiated in the brain. But he famously observes that while we must use our sympathetic imagination to conceive of mental states, with their qualities of sensation and emotion, we must use our very different perceptual imagination to conceive of the brain, with its qualities of viscosity and lumpishness. And this qualitative difference is exactly why “we can imagine,” Nagel says, an object like the brain without having to imagine, as well, those properties—the mental states—that it might instantiate (see also Worley 2003:22); in other words, we can imagine zombies. Suppose that we can, then, explain how the unique instantiatonal relationship between brain and mental states requires no evident similarity in qualities—perhaps because we access them through entirely different imaginative faculties—and in addition, that instantiation captures their object-property difference in ontologies. Then, it would seem, instantiation gives us the beginnings of conceptual purchase on the hard problem.

Kripke, too, frames his discussion in terms of “instantiation” (Kripke 1980:135, 136). Doing so enables him then to distinguish a physical property like heat, instantiated in objects like moving molecules, from a mental property like pain, instantiated in an object like the brain. Heat is necessarily instantiated in moving molecules. And so those moving molecules cannot be imagined without their instantiating the property of heat. True, Kripke (1980:132) adds, they can be imagined without their instantiating the “property that produces. . .sensations” of heat “in people.” But the property of pain just is a sensation. And so firing C-fibers can easily be imagined not to instantiate that property. It’s possible, then, to use instantiational conceptualization to conceive of a zombie. That’s precisely because—in the case of mental properties and the brain, but unlike the case of heat and moving molecules—we needn’t imagine any of the qualities of the property when we imagine the totally different qualities of the instantiating object.

There’s a parallel here between Olson’s compositional discussion of zombies and the Nagel/Kripke instantiational approach. Olson needs to supplement his compositional conceptualization with an additional “metaphysical principle” that accommodates the ontological difference between mental state and brain in a way that compositional conceptualization alone doesn’t. In the same way, Nagel and Kripke supplement instantiational conceptualization with arguments—about the unique ways in which we conceive of the instantiation of mental as opposed to physical properties—that accommodate the qualitative differences between mental states and brain in a way that instantiational conceptualization alone does not. That very difference in qualities, however, is precisely what renders instantiational conceptualization—requiring as it does some qualitative similarity between object and property—inadequate to the task. If in accounting for the ontological difference between mental states and brains, instantiation at the same time relies on their qualitative dissimilarity to do so, then the hard problem simply reappears. What could the connection be between such vastly different physical and mental qualities, between the sponginess or viscousness of the instantiating object, and the pain or joy of the instantiated mental state? Instantiational conceptualization cannot account for qualitative differences—although, of course, compositional conceptualization can, and indeed is used to do so.

By tracing the contours of the debates over multiple realizability and zombies, we can see that what we might need—what we might be grappling for—is a certain kind of conceptual breakthrough. Perhaps what we require, or have yet to devise, is a hybrid of those two basic classes of concepts, compositional—with its capacity for explicating physical-mental differences in qualities (but not ontologies) —and instantiational, with its capacity for explicating physical-mental differences in ontologies (but not qualities). Perhaps we might then take the next step in conceiving of how mental states float so free of their physical substrates, differing in both their qualities and their ontologies, and yet so thoroughly depending on them.

3.1 Mental Causation and Panpsychism

Turn now from multiple realizability and zombies to two other contemporary discussions the hard problem provokes: the ones surrounding mental causation

and panpsychism. Here, the conceptual quest differs in a key way from the one sparked by multiple realizability and zombies. Instead of trying to account for how the mental and the physical on which it depends can differ so profoundly in both their qualities and their ontologies, the challenge is the opposite: to show how the mental and the physical, while remaining distinct, might somehow share some of the same qualities and the same ontologies. After all, if a mental state like a headache in some sense shares in the ontology and the neuromechanical qualities of its physical substrates, then it's conceivable as to how it might participate in causing me to take a physical action in the material world, like reaching for analgesics. And if physical particles in my brain in some sense share in the ontology and phenomenal qualities of my mental states, as panpsychists say, then it's possible to see how they can form and embody my conscious mind and thus help dissolve the hard problem.

Here again, in the attempts to conceptualize both mental causation and panpsychism, the two schemata, compositional and instantiational, form a leitmotif. Composition does so because it requires, and hence can more readily account for, shared mental-physical ontology. And instantiation does so because it requires, and can more readily account for, shared mental-physical qualities.

4. Mental Causation, Instantiational Conceptualization, and Compositional Conceptualization

Instantiational concepts, to begin with, have assisted in bolstering the conceivability of mental causation. For Shoemaker (2007:2), we can say that the brain's physical substrate instantiates mental states like pain because we can, in fact, identify a particular quality of the mental state in the physical substrate. Specifically, the quality that the brain's instantiating physical substrate and the instantiated mental state share is the “conditional power” to cause me to reach for the pill bottle, the word “conditional” signaling that the power requires—if I actually am to reach for it—other events to occur, such as my arm moving.

So just as instantiated properties and their instantiating objects do, the mental state of pain and its substrates in the brain share a quality that allows us to say that the mental state is instantiated in the brain. Here, what they share is the conditional power to cause me to reach for the aspirin. It is that conditional power, a quality that belongs every bit as much to the instantiated mental state as to its instantiating physical substrate in the brain, which causes my physical action. The way in which the brain instantiates mental states thus helps explain how those mental states, along with their physical substrates in the brain, can cause physical events.

In Shoemaker's terminology, the brain—firing C fibers—“instantiates” the mental state of pain because the conditional power to cause me to reach for the aspirin, qua physical, “realizes” the conditional power to cause me to reach for the aspirin, qua mental (Shoemaker 2007:14, 21, 27, 28–9). And how does that realization take place? Here Shoemaker combines his instantiational conceptualization—on which the instantiating brain and the instantiated mental state of pain share the same quality, the same conditional causal power—with

compositional conceptualization, on which the conditional power qua physical is a whole, which realizes the conditional power qua mental as a part. The conditional power to cause me to reach for the aspirin, qua physical, can also cause a P1 reading on a cerebroscope; the conditional power to cause me to reach for the aspirin, qua mental, cannot. This makes the conditional power qua mental a part, and the conditional power qua physical the whole: after all, wholes are larger than and inclusive of their parts.

By realizing the conditional power qua mental as a part, the conditional power qua physical as a whole endows it with its ontological status as a physical property, enabling it to participate in causing physical events. And, of course, since parts and their wholes don't compete as casual forces, both can cause the same event. As Wilson (2011:131) notes, "proper parts are distinct from and yet in a sense nothing over and above wholes, and may be efficacious without inducing overdetermination, as when both I, and my eye, cause a wink" (see also Clapp 2001; Yablo 1992; Pineda and Vicente 2017:100).

Certainly, mental states and the brain are in many ways profoundly qualitatively different, the headache pain of the one nowhere evident in the dull grey matter of the other. But Shoemaker shows that they actually do share a quality—they each possess the conditional power to cause me to reach for the aspirin—that allows us to see how the brain instantiates mental states. Mental states and the brain also belong to different ontological categories, properties and objects. But Shoemaker shows how—if the conditional power qua mental can in some sense be understood as a part that helps compose the conditional power qua physical as a whole—part and whole can then share the same "ontological status" as physical properties. As Shoemaker (2007:114) says, "the ontological status" of the causal power qua mental just "is that of macroscopic [causal powers qua] physical." That allows the mental state of pain to participate, along with and just as much as the brain, in causing an event in the physical world.

It is revealing, then, that Shoemaker relies on both compositional and instantiational conceptualization. Each is needed. The instantiational is needed to suggest that the mental and the physical do share a common quality, the same conditional causal power. And the compositional is needed to suggest that they also share a common ontology, such that they can actually both in some way participate in causing the same physical event. My question here is not whether Shoemaker's foundational theory about mental causation succeeds. It's whether insofar as it runs into difficulties, at least some of them might have to do with the conceptual problems we face in combining instantiational and compositional frames. Mental causation, in Shoemaker's rendering, requires both.

Notably, a key line of criticism directed at Shoemaker's view does focus on whether the mental and the physical can simultaneously have both a compositional and an instantiational relationship. As Audi (2012:669) for example notes in his discussion of Shoemaker, "it is not plausible to regard wholes as realizing their proper parts, or in any other way as ontically prior to them. . . The fact that the chair exists, e.g., does not account for the fact that the legs, seat, and back exist. It is the other way around." In other words, since, for Shoemaker, the brain instantiates the mental state of pain just because one of its

particular powers qua physical “realizes” the same particular power qua mental, the power qua mental could not then at the same time be part of the power qua physical. Or as Audi also puts his point, if, as it is on Shoemaker’s view, the power qua physical is the whole, then it is not “ontically prior to what [it is supposed to] realize”, namely the power qua mental, and so in fact cannot realize it. For Audi, this means that views like Shoemaker’s fail “to accommodate a basic feature of” what’s required for the brain to instantiate mental states.

More generally: The physical cannot at one and the same time be a whole composed in part by the mental, and be an instantiator of the mental. On the one hand, if the causal power qua mental relates to the causal power qua physical as part to whole, then while we can conceive of them as sharing an ontology, the physical could not then be readily constructable as realizer of the mental. That, however, is what is required to create conceptual space for the idea that the mental state of pain is instantiated in the brain, displaying at least some of the same qualities (the conditional power to cause me to reach for the analgesic) along with it. On the other hand, if the brain’s physical substrate does instantiate the mental state, the two can certainly share a quality, namely the causal power. But then the causal power qua physical cannot easily be constructable as a whole sharing an ontology with the causal power qua mental as a part, which is what allows for them both to cause me to reach for the analgesic.

I am not taking a position on how successful Audi’s criticism is (consider a view on which wholes are not dependent on their parts, but vice versa; see e.g., Shani 2015). Several others, though, have advanced similar critiques of Shoemaker. Kim (2010:111–12), for example, discusses the tension involved when a mental state both helps “compose” and is “instantiate[d]” by the physical in Shoemaker’s philosophy (see also Pineda and Vicente 2017, and Morris 2011). My point is that the emergence of this vein of criticism furnishes an additional indicator, this time in the mental-causation debate, of the difficulty involved in drawing on the compositional and the instantiational—combined with the apparent need to do so.

5. Panpsychism, Compositional Conceptualization, and Instantiational Conceptualization

Turn now to the use of compositional and instantiational conceptualization in panpsychism. Suppose that each of the innumerable microparticles in my brain were itself conscious, even if at an extremely elementary level. Then my conscious mind might be composed from the “consciousness of its most fundamental parts” (Goff, Seager, and Allen-Hermanson 2022:Sec. 3.1).

In other words, just as we can imagine how some ontological objects can be parts that compose another ontological object as a whole—as my car’s engine, tailpipe, tires, etc., compose the car itself—so it’s possible to conceive of how all those ontological properties, those mini-mental states in my brain, could be parts that together (perhaps through “phenomenal bonding” or “mental chemistry”) compose another ontological property, namely my full-blown mental state, as a whole. After all, their ontology is the same. As Coleman (2014:21) says, “it has been panpsychism’s *raison d’être* to preserve ontological continuity between the

world's lower ontological levels and the conscious macro-mentality arising from these levels." Compositional whole-part conceptualization helps panpsychists to do so.

Here, the micro-mental states—the parts—might be ontological properties every bit as much as is the conscious mental state, the whole, that they compose. But then, so a prominent line of criticism has it, with their dim, primitive, and opaque impressions and apprehensions, those micro-mental states wouldn't necessarily display the same qualities—the vibrant emotions, crisp cognitions, and spectacular sensations—which that macro-mental state as a whole displays. The qualities of parts and whole here differ dramatically, as is often the case with composition (O'Connor and Wong 2005:665). As Papineau (2020:29) says, even if "microscopic parts. . .are credited with some conscious nature"—rudimentary sensations, say, or primitive apprehensions—these will "presumably be different in kind from the conscious nature of the wholes they compose," with their exquisite sensations and magnificent cognitions (see also Coleman 2014:38; Roelofs 2020:246; Goff 2017:182). And so "why," Papineau asks, "is the relation between the conscious parts and the differently conscious wholes [they compose] any less mysterious than the supposedly puzzling emergence of conscious wholes from non-conscious parts?" In other words, how is the problem of panpsychism any different from the hard problem itself?

For panpsychism, compositional frames thus run into resistance. And a principal reason is that while they can account for the necessary similarities in ontology between microconsciousnesses and macroconsciousness, they can't, for critics, bridge the consequent difference in their qualities. And so, again, panpsychist philosophers have turned to instantiation concepts, which do have the capacity to remedy this failure of compositional conceptualization.

Coleman's use of instantiation conceptualization illustrates its capacity to bridge the qualitative difference between micro and macro consciousness. A macro-mental state such as "feeling cold, tired and smelling roast beef," Coleman (2014:23–4) says, exists when "one set of ultimates" at the micro level "instantiate[es] phenomenal coldness, another set instantiate[es] phenomenal tiredness, and a further set instantiate[es] roast beef smell." Instantiation eliminates the qualitative difference between micro and macro, because different "sets" of "ultimates" at the micro level (for most panpsychists, those ultimates are objects such as the brain's microparticles) directly instantiate robust macro properties like tiredness, coldness, and roast-beef smell.

Still, it has been hard for many to see how each of those instantiated ontological properties, e.g., the macro-mental state of tiredness, or coldness, could cast off the object ontology of its particular set of instantiating microparticles, and coalesce with one another into a unified macro-mental state, to make for a combined consciousness of cold and tiredness. Here, critics often cite William James' famous formulation. A hundred people standing next to each other might all instantiate different fully conscious states. But their discrete, ontological embodiments as objects—their distinct spatio-temporal locations in the material world, as Coleman (2014:30) says—would bar them from instantiating a unified ontological property such as a combined mental state. The ontological difference between

properties and objects means that each conscious state, each property, would still remain sequestered by the physical barriers of its instantiating object, “shut in its own skin” as James says (see relatedly, Lewtas 2018, 399), and unable to combine.

Suppose that the “ultimates” are not microparticles but microsubjects. Then as Goff (2017:183) says:

[w]e might imagine. . .that some. . .micro-subjects underlying relevant brain functioning have an experience as of staring at grass, some have an experience as of freshly cut grass smell, and some have an experience as of the sound of a distant lawnmower. Still, all of these micro-level goings on seem consistent with the absence of some subject having a *unified experience* as of seeing and smelling grass while hearing a distant lawnmower. No matter what weird and wacky conscious states we attribute to your micro-level bits, it seems conceivable that those micro-subjects exist in the absence of some further subject at the macro-level.

Each group of “micro-subjects,” in instantiating a particular macro-experience—whether the smell of grass, the sight of grass, or the sound of a lawnmower—allows for a panpsychist qualitative sameness of consciousnesses at the micro and macro levels. But micro-subjects and macro-experiences aren’t ontologically the same: micro-subjects are ontological objects, and so can’t combine so as to instantiate the property of a single “unified experience” at the macro level.

Suppose that we conceive of the “sets of ultimates” that instantiate macro-experiences not as micro-particles or micro-subjects but, as Coleman suggests, simply as micro-experiences or micro-phenomena of some sort. In that case, it would seem that not only would the “qualities,” as Chalmers (2016:205) puts it—coldness, fatigue, and smelling roast beef—be the same at both micro and macro levels, so would the ontology: the property ontology of experiences or phenomena. Since the micro-experiences or micro-phenomena are properties, not objects, they aren’t segregated physically, shut in their own skin, and so should be able to combine to instantiate a unified macro-mental state. And yet even here, for critics, instantiatational conceptualization cannot finally evade the problem of ontological difference. As Roelofs (2020:253) notes:

the brain seems to retain all the physical features of its parts, including the incredibly specific. . .distribution of features across microscopically small locations. . .Perhaps the experiential parts of the brain fuse. . .but the physical parts do not correspondingly fuse into a single physical entity. . .But then it remains unclear how this would work: the difference in ontology seems to persist.

Separate macro-experiences like coldness, or fatigue, or roast-beef smell—which are all ontological properties—might be instantiated in different “experiential parts of the brain” at the micro level. But the corresponding physical parts of the brain—

those ontological objects—cannot “fuse,” putting in question the possibility of a unified macroexperience. “The difference in ontology,” as Roelofs says, “seems to persist.”

Instantiational conceptualization thus accounts for how the brain’s micro-level entities could share qualities with its macro mental states, boosting panpsychism, but not—for its critics—the same ontology, which panpsychism also requires. Meanwhile compositional conceptualization, as I have suggested, might allow for micro and macro mental states to share the same ontology, but at the cost of their being seen to have very different qualities (see also some of the discussion between Bohn (2019) and O’Conaill (2022)). The debate over panpsychism is of course sprawling and complex. My concern has been to suggest how a basic issue it faces in generating an answer to the hard problem—the complementary inadequacies of both compositional and instantiational conceptualization—resembles those faced in those other, equally sprawling debates: over multiple realizability, zombies, and mental causation.

6. Conclusion

My aim here has not been to plumb the full depths of the different arguments I have examined, but to abstract a broad commonality between them. It’s significant that across all four debates discussed here, philosophers have deployed the same two conceptual schemes, composition and instantiation, to grapple with the hard problem. And in all four, they have laid bare much the same complementary difficulties faced by each scheme.

In the debates over multiple realizability and zombies, composition accounts for the necessary difference in qualities between the mental and the physical but not—or not yet—for the necessary difference in ontology. Instantiation does the reverse: It readily accounts for the necessary mental-physical differences in ontology but not for the difference in their qualities. Meanwhile, in the debates over mental causation and panpsychism, compositional conceptualization can account for the necessary mental-physical similarities in ontology, but not—by itself—for the necessary similarities in mental and physical qualities. Instantiation, for its part, accounts for the necessary similarities in qualities, but not—by itself—for the necessary similarities in ontology.

My aim has not been to take sides. Instead, it has been to pull back and gain a broad-gauge view of the four debates together, lending some clarity to a common conceptual challenge they have collectively bumped up against. What all of these efforts do is give us a glimmer of at least one central set of conceptual limitations we must break through even if there might be others.

Any such breakthrough would almost certainly be insufficient on its own to resolve the hard problem, but it could well be necessary. The really hard question, then, might be whether such a breakthrough lies beyond our cognitive capacities. If it doesn’t lie beyond those capacities, then overcoming the difficulties posed by compositional and instantiational conceptualization would seem to point in a direction where we might begin to gain a new purchase on the mystery of consciousness. If overcoming those difficulties in fact does lie beyond our

conceptual capabilities, then we will at least gain some purchase on how our mental apparatus might prevent us from resolving the hard problem.

Acknowledgements. I am grateful to Paul Thompson, Jessica Wilson, and the anonymous reviewers for *JAPA* for their helpful comments on earlier versions of this article.

ANDREW STARK

UNIVERSITY OF TORONTO, TORONTO, CANADA

stark@rotman.utoronto.ca

References

- Aizawa, Kenneth and Carl Gillett. 2009. “The (Multiple) Realization of Psychological and other Properties in the Sciences.” *Mind and Language* 24: 181–208.
- Antony, Louise M. 1999. “Multiple Realizability, Projectibility, and the Reality of Mental Properties.” *Philosophical Topics* 26: 1–24.
- Audi, Paul. 2012. “Properties, Powers, and the Subset Account of Realization.” *Philosophy and Phenomenological Research* 84: 654–674.
- Aydede, Murat and Güven Güzeldere. 2005. “Cognitive Architecture, Concepts, and Introspection: An Information–Theoretic Solution to the Problem of Phenomenal Consciousness.” *Nous* 39: 197–255.
- Bennett, Karen. 2017. *Making Things Up*. New York: Oxford.
- Bohn, Einar Duenger. 2019. “Panpsychism, The Combination Problem, and Plural Collective Properties.” *Australasian Journal of Philosophy*, 97 (2): 383–394.
- Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford.
- Chalmers, David. 2008. “Strong and Weak Emergence.” In Philip Clayton and Paul Davies (eds.), *The Re-Emergence of Emergence*. New York: Oxford, 244–254.
- Chalmers, David. 2016. “The Combination Problem for Panpsychism.” In Godehard Bruntrup and Ludwig Jaskolla (eds.), *Panpsychism: Contemporary Perspectives*. New York: Oxford, 179–214.
- Clapp, Lenny. 2001. “Disjunctive Properties: Multiple Realizations.” *Journal of Philosophy* 98 (3): 111–136.
- Coleman, Sam. 2014. “The Real Combination Problem: Panpsychism, Micro–Subjects, and Emergence.” *Erkenntnis* 79: 19–44.
- Doggett, Tyler and Daniel Stoljar. 2010. “Does Nagel’s Footnote Eleven Solve the Mind–Body Problem?” *Philosophical Issues* 20: 125–143.
- Fisher, A.R.J. 2018. “Structural Universals.” *Philosophy Compass* 13: 1–13.
- Gillett, Carl. 2007a. “The metaphysics of mechanisms and the challenge of the new reductionism,” in Maurice Schouten and Huib Looren de Jong (eds.), *The matter of the mind*. Oxford: Blackwell.
- Gillett, Carl. 2007b. “Understanding the New Reductionism: The Metaphysics of Science and Compositional Reduction.” *Journal of Philosophy* 104: 193–216.
- Gillett, Carl. 2016. *Reduction and Emergence in Science and Philosophy*. New York: Cambridge.
- Goff, Philip. 2017. *Consciousness and Fundamental Reality*. New York: Oxford.
- Goff, Philip, William Seager, and Sean Allen-Hermanson, “Panpsychism”, *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2022/entries/panpsychism/>
- Kim, Jaegwon. 2010. “Thoughts on Sydney Shoemaker’s Physical Realization.” *Philosophical Studies* 148: 101–112.
- Kripke, Saul A. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.
- Kroedel, Thomas. 2020. *Mental Causation: A Counterfactual Theory*. New York: Cambridge University Press.
- Lewtas, Pat. 2018. “Panpsychism, Emergentism, and the Metaphysics of Causation.” *Pacific Philosophical Quarterly* 99: 392–416.

- Macdonald, Cynthia. 2005. *The Varieties of Things*. Oxford: Blackwell.
- Macdonald, Cynthia and Graham Macdonald. 1986. "Mental Causes and Explanation of Action." *Philosophical Quarterly* 36: 145–158.
- McGinn, Colin. 1989. "Can We Solve the Mind-Body Problem?" *Mind* 98: 349–366.
- Morris, Kevin. 2011. "Subset Realization, Parthood, and Causal Overdetermination." *Pacific Philosophical Quarterly* 92: 363–379.
- Moser, Paul K. and J.D. Trout. 1995. "Physicalism, Supervenience, and Dependence," in Elias E. Savellos and Umit D. Yalcin (eds.), *Supervenience: New Essays*. New York: Cambridge, 187–217.
- Nagel, Thomas. 1974. "What Is It Like to Be a Bat?" *Philosophical Review*, 83 (4): 435–450.
- Nagel, Thomas. 2012. *Mind and Cosmos*. New York: Oxford.
- O'Conaill, Donnchadh. 2022. "Panpsychism, Emergence, and Pluralities: Reply to Bohn." *Australasian Journal of Philosophy*, 100 (2): 419–424.
- O'Connor, Timothy and Hong Yu Wong. 2005. "The Metaphysics of Emergence" *Nous* 39: 658–678.
- Oliver, Alex. 1996. "Metaphysics of Properties." *Mind*, 105: 1–80.
- Olson, Eric. 2018. "The Zombies Among Us." *Nous* 52: 216–226.
- Papineau, David. 2020. "The Problem of Consciousness." In Uriah Kriegel, ed., *The Oxford Handbook of the Philosophy of Consciousness*. New York: Oxford, 14–36.
- Pineda, David and Agustin Vicente. 2017. "Shoemaker's Analysis of Realization: A Review." *Philosophy and Phenomenological Research* 94: 97–120.
- Polger, Thomas W. and Lawrence A. Shapiro. 2016. *The Multiple Realization Book*. New York: Oxford.
- Rea, Michael. 1995. "The Problem of Material Constitution." *Philosophical Review* 104 (4): 525–552.
- Roelofs, Luke. 2020. "Can We Sum Subjects?" in William Seager, ed., *Routledge Handbook of Panpsychism*. London: Routledge, 245–258.
- Rudder Baker, Lynne. 2019. "Intentionality and Emergence" in Sophie Gibb *et al.* eds., *Routledge Handbook of Emergence*. London: Routledge, 206–214.
- Shani, Itay. 2015. "Cosmopsychism: A Holistic Approach to the Metaphysics of Experience." *Philosophical Papers*, 44 (3): 389–437.
- Shoemaker, Sydney. 2007. *Physical Realization*. Oxford: Clarendon.
- Silberstein, Michael. 1998. "Emergence and the mind–body problem" *Journal of Consciousness Studies* 5: 464–482.
- Williams, J.R.G. 2007. "The possibility of onion worlds: Rebutting an argument for structural universals." *Australasian Journal of Philosophy*, 85 (2): 193–203.
- Wilson, Jessica. 2011. "Nonreductive realization and the Power-Based Subset Strategy." *The Monist*, 94: 121–154.
- Wilson, Jessica. 2021. *Metaphysical Emergence*. Oxford: Oxford University Press.
- Worley, Sara. 2003. "Conceivability, possibility and physicalism." *Analysis* 63: 15–23.
- Yablo, Stephen. 1992. "Mental Causation." *Philosophical Review*, 101 (2): 245–280.