



Replies to Commentators

Michael Pelczar¹

Received: 28 February 2025 / Revised: 2 April 2025 / Accepted: 4 April 2025
© The Author(s), under exclusive licence to Springer Nature B.V. 2025

1 Introduction

An author couldn't ask for more intelligent, incisive, or constructive commentary than the contributors to this forum have provided. My only complaint is that they've provided more than I can do full justice to within manageable bounds. Fortunately, there are recurring themes in their contributions, with different commentators often honing in on the same issue. To keep an already lengthy discussion to reasonable size, I've decided to take advantage of this convergence by organizing my replies thematically, rather than by commentator (though I'll try to signal who I'm responding to at any given stage). I won't claim to have addressed all the points the commentators raise, but I've tried to address the most important ones.

My replies fall into roughly two categories. The first includes my replies to internal threats raised against my theory (ostensible shortcomings intrinsic to the theory itself); these include threats related to the theory's reliance on the possibility of minds very unlike our own (§2) and its egocentric slant (§3). The second category includes my replies to external threats (respects in which my theory is alleged to compare unfavorably to competing theories); these include threats from various forms of metaphysical realism (§5) and panpsychism (§6). §4 addresses a combination of internal and external threats, related to phenomenalism's alleged disadvantages to a causal theory of physical objects, and my effort to overcome them with an account of ideal time and space.¹

¹ Very roughly, I'm responding to DeRosset primarily in §2 and §3, to Fumerton primarily in §2, §3, and §4, to Hedda Hassel Mørch primarily in §6, and to Smithson primarily in §2 and §5.

✉ Michael Pelczar
pelczar.philosophy@gmail.com

¹ Department of Philosophy, National University of Singapore, Singapore, Singapore

2 Possible Minds

In my view, the physical facts of our world reduce to facts about what sorts of experience are likely to occur, given the occurrence of other sorts of experience. Right now I'm seeing a stack of books. The visual experiences I'm having belong to a bigger multi-modal conscious state (that also includes auditory experiences, tactile experiences, etc.), which itself belongs to a complete stream of consciousness (currently in progress), which in turn belongs to a whole mental life (also in progress). Though I'm the only one seeing these books right now, there are lots of possible scenarios where other people see them too. In all of them, there are mental lives that include bookish experiences like the ones I'm having now. Each of these possible mental lives is such that the probability of its including such experiences, conditional on its having all the other conscious features it has, is high (or at least, greater than chance). Furthermore, each of the possible mental lives is such that the probability of its bookish experiences having certain conscious properties given that the other mental lives' bookish experiences have certain conscious properties—e.g., various phenomenal shapes and sizes, or the property of undergoing various phenomenal changes—is also high (or at least greater than chance).²

According to the theory I favor, the existence of the stack of books is no more, and no less, than the existence of these and related probabilities: probabilities for certain purely experiential states of affairs to exist conditional on the existence of other purely experiential states of affairs. I call such probabilities “phenomenal probabilities;” they are my take on Mill's permanent possibilities of sensation. In my view, physical things are phenomenal probabilities.³

Several commentators suggest that in building my theory, I fail to take into account the very different ways in which different minds might engage with the physical world, to the theory's detriment.⁴

DeRosset frames the criticism as an apparent conflict between two claims, both of which I need to be true. First claim: the class of possible minds is inclusive enough to let me construct all the physical phenomena we have reason to believe exist, including

² Though I don't discuss these in the book, there are also relevant higher-order probabilities, such as the probability that each mental life includes the relevant bookish experiences, conditional on (1) its having all the other conscious features it has and (2) each of the other possible lives being such that the probability of its occurring with the relevant bookish experiences given that it occurs with all its other conscious features is greater than chance.

³ I never go so far as to offer a full translation of a physical statement in terms of phenomenal probabilities; as DeRosset surmises, my goal is to provide a proof of concept of the basic phenomenalist idea, rather than detailed analyses of specific physical facts. In this respect, my book is par for the course. Anyone looking for granular analyses of specific instances of a reductionist metaphysician's target phenomenon is in for a disappointment. This is so whether the target is physical reality or, e.g., mental reality.

⁴ Fumerton puts the challenge in terms of a “perceptual relativity” by which “what sequences of experiences we have in any given situation is relative to external conditions *and* characteristics of the relevant *perceiver*.” Smithson puts it in terms of the possibility for “idiosyncratic experiences.” DeRosset puts it in terms of the demarcation problem outlined below.

phenomena, like quarks, that the laws of nature prevent anyone from actually perceiving. Second claim: the class of possible minds is restrictive enough to prevent me from inadvertently constructing physical things we know do not exist, like unicorns. The criticism is that no class of possible minds can make both these claims true. In short: I can't have quarks without unicorns, and I can't renounce unicorns without renouncing quarks.

Let's put the challenge as an argument. Taking my cue from DeRosset, I'll call it the *Demarcation Argument*:

1. There are nomically imperceptible physical phenomena.
2. If there are nomically imperceptible physical phenomena, then phenomenalists have to define physical things partly in terms of possible minds radically unlike our minds.
3. So, phenomenalists have to define physical things partly in terms of possible minds radically unlike our minds. (1, 2)
4. If phenomenalists have to define physical things partly in terms of possible minds radically unlike our minds, they have to accept the possibility of oddball minds.
5. So, phenomenalists must accept the possibility of oddball minds. (3, 4)
6. If oddball minds are possible, then phenomenalism implies the existence of physical phenomena that do not, in fact, exist.
7. So, phenomenalists must accept the existence of physical phenomena that do not, in fact, exist. (5, 6)

Premise 1: Nominally imperceptible physical phenomena are physical phenomena that could not be perceived without a violation of natural law. Whether there are such phenomena is disputable (scientific antirealists dispute it). However, I don't want to deny that there are. I mean my theory to be compatible with the existence of physical things like quarks. So I accept this premise.⁵

Premise 2: I accept this too. As DeRosset points out, any minds capable of perceiving things that the laws of nature prevent us from perceiving are bound to be very unlike our minds. Yet, it is only out of possibilities for the sort of experiences such minds can have that I can construct things like quarks. The alternative is to try to reduce things like quarks to possibilities for the kinds of human experiences that justify human observers in believing there are quarks. But, for reasons I give in the book (pp. 134–35), the prospects for this seem poor. If a quark is a possibility of sensation, it's a possibility for sensations occurring in minds radically unlike ours.⁶

⁵ I also mean my theory to be compatible with claim that the physical contents of our world could have existed even if *all* experience had been nomically impossible, as in a zombie world. In a zombie version of our world, we can't refer to physical things (or, perhaps, to anything), but physical things still exist. (More on this in §3.)

⁶ That said, I don't think it's correct to suggest (as Hedda Hassel Mørch does) that phenomenalism "posits more types of experiences" than other theories. Phenomenalism posits no more experience than we have independent reason to believe exists; in this respect, it's more parsimonious than panpsychism, which posits an astronomical number of inscrutable experiences to play the role of fundamental physical phenomena. (Phenomenalism presupposes a great number of *logically possible* experiences, but no more than other theories.)

Premise 4: DeRosset describes a population of disembodied minds who are disposed to perceive a certain region of empty space as containing a cantaloupe. In other respects, they have experiences much like our own. It's just that when they direct their attention to this particular region of space, they tend to have the sort of experiences we have when perceiving a cantaloupe. Their experiences relate to one another in just the way that different human beings' experiences relate when the humans perceive the same cantaloupe; in my jargon, their cantaloupish experiences constitute a "thinglike group."⁷

The minds DeRosset describes are an example of what I'll call "oddball minds." The defining characteristic of an oddball mind is that (1) it belongs to a population of minds (or possible minds) whose actual and possible experiences include some that have exactly the same qualities and stand to each other in exactly the same dependence relations as the experiences of human beings who all perceive some physical thing, despite the fact that (2) there is no actual physical thing answering to those experiences.

There's no limit to the oddballs we can dream up. For any cantaloupe-shaped region of space, we can imagine a population of oddballs who have thinglike experiences suggestive of a cantaloupe occupying that region. We can imagine a population of oddballs who are disposed to have thinglike experiences of a unicorn grazing by the seventh hole of Augusta National golf course, where there's really nothing but an empty putting green; etc.

DeRosset, Fumerton, and Smithson all suggest that I might want to deny Premise 4. To do so, I would have to articulate a sense of "possible" in which minds with experiences sufficient for a phenomenalist construction of quarks are possible, but minds with experiences sufficient for a phenomenalist construction of unicorns are not.

This isn't the approach I want to take.

It's hard to see how to limit the range of relevantly possible minds other than by reference to similarities to our own minds (or, as Fumerton puts it, by "restrict[ing] the relevant conditional probabilities to those involving human experience"). But that leads straight back to the quark problem. I might *define* "relevantly possible minds" as those whose associated possibilities combine with my theory to entail the existence of all and only the physical things that actually exist. But that would be blatantly circular, given my aims. If I want to define "relevantly possible mind" non-circularly, and without unduly restricting the range of possible minds to those similar to ours, then how? I doubt it can be done.

So, in good Millian fashion, I counsel toleration: in the phenomenalist community of possible minds, all are welcome, including oddballs! I accept Premise 4.

Premise 6: This is where the Demarcation Argument goes wrong. To see why, it'll help to step back and look at the argument as a challenge to a broader class of theories

⁷ In a thinglike group, experiences occur with the qualities and relations that typically characterize the experiences that occur in different people when they perceive the same thing.

that includes not just phenomenalism but causal theories of the physical like noumenalism, according to which the physical facts of our world reduce to facts about the experience-causing powers of the entity or entities—further nature(s) unknown—that have experience-causing powers (the so-called “noumena”). As Fumerton points out, the threat to phenomenalism that the Demarcation Argument crystallizes is also a threat to causal theories.⁸

Take the Demarcation Argument and replace all the occurrences of “phenomenalism” and “phenomenalists” with “noumenalism” and “noumenalists.” Noumenalists should accept premises 1, 2, and 4 of the resulting argument. If they want to allow for things like quarks, they have to define physical reality partly in terms of noumenal powers to cause experiences in minds radically unlike ours; having let such minds through the door, they have no principled basis for excluding the possibility of all manner of oddball minds, including populations of minds that the noumena are apt to give thinglike experiences of unicorns and interstellar cantaloupes. If the existence of noumenal powers to cause thinglike experiences of unicorns is enough to commit noumenalists to the existence of unicorns, noumenalists are in trouble (at least, if they’re realists about quarks).

But the existence of noumenal powers to cause thinglike experiences of unicorns is *not* enough for the existence of unicorns, by the standards of a properly formulated noumenalism. For there to be unicorns, the noumena must have the power to cause unicornish experiences that (1) occur in thinglike groups, and (2) cohere with all the other experiences that occur in the nearest ideal world. The oddballs’ unicornish experiences satisfy (1), but not (2).⁹

An “ideal world” is a hypothetical totality of experiences that includes all the experiences you’d expect to exist if there were sentient beings who collectively perceived every physical feature of some universe. I call such beings “ideal observers,” and the totality of their experiences a “worldlike totality.” Imagine some physical spacetime—the one we actually inhabit or some other—partitioned into regions, each of which has its physical contents monitored by a team of disembodied sentiences. Now bracket spacetime and its physical contents, and just consider the manifold of experiences that occur in those sentiences with all their counterfactual or probabilistic interde-

⁸ Here I might as well confess that, like Fumerton, I find building physical objects with our understanding of causation to be in a sense more intuitive than constructing physical objects out of counterfactuals or conditional probabilities. However, I don’t think this weighs in favor of causal theories or against phenomenalism. I think it just reflects the fact that my hominid brain has evolved to think in terms of cause-and-effect more readily than in terms of counterfactual dependence or conditional probability.

⁹ To relate this back to DeRosset’s discussion, my view does indeed yield the truth of DeRosset’s proposition (5) (given at the end of Section 4 of his contribution), but it doesn’t imply the truth of the proposition you get if you add to (5) the following clause: “that cohered with the other experiences in that ideal world.” (The solution to the demarcation problem that I propose here is a third strategy, additional to the two that DeRosset considers.)

dependencies. By an ideal world, I mean a totality of experiences that includes such a manifold.¹⁰

In effect, an ideal world is the kind of thing Leibniz thinks the *actual* world is. *Pace* Leibniz, our experiences do not, I take it, actually belong to an ideal world, but we can still speak of the ideal world that they *would* belong to *if* they belonged to one, or the ideal world that our experiences have the highest probability of belonging to, conditional on their belonging to an ideal world.

The *nearest* ideal world is the one most likely to exist, conditional on there being an ideal world that includes all of our actual experiences. It's the logically possible ideal world W such that (1) W includes all our actual experiences, and, (2) for all logically possible ideal worlds ω , if $\omega \neq W$, then $\Pr(W \text{ exists} \mid \text{our actual experiences belong to an ideal world}) > \Pr(\omega \text{ exists} \mid \text{our actual experiences belong to an ideal world})$. Roughly, the nearest ideal world is the least improbable scenario where Leibniz's metaphysics is true.

An ideal world includes many (perhaps infinitely many) groups of experiences that relate to each other in the ways you'd expect of the experiences occurring in different conscious beings perceiving the same physical things at the same time. Since the nearest ideal world includes all our actual experiences, it also includes experiences that don't belong to such groups, such as our non-perceptual experiences (like feelings of giddiness or depression), quasi-perceptual experiences (like after-images), and the experiences we have in dreams and hallucinations.

Not even all thinglike groups of experiences that occur in the nearest ideal world correspond to real physical things. That honor goes only to thinglike groups whose constituent experiences cohere with the totality of experiences with which they co-exist (as features of the same logically possible world).

Here I would like to give a sharp, compact definition of "cohere." Sadly, I cannot. What I can do is give examples of experiences that do *not* cohere with others, and ostensively define coherence as the relation that these experiences saliently fail to bear to other experiences.¹¹

There's a way that the experiences you're having right now relate to all the other experiences you've had, and that the experiences you've had in dreams do *not* relate to the other experiences you've had. There's a way your current experiences relate to

¹⁰ This way of introducing the concept of an ideal world (as including all the experience that would occur if all the physical features of some universe were perceived) doesn't make an analysis of physical things in terms of ideal worlds circular. I refer to physical worlds in order to give you a certain concept: the concept of a state of affairs in which there are experiences with certain phenomenal qualities and standing in certain relations of interdependence; once you've got the concept, you can entertain such a state of affairs without assuming the existence of an underlying physical reality to support or explain it. Analogously, if I give you the concept of fool's gold by telling you that it's a compound of iron and sulfur that superficially resembles gold, you can go on to entertain a world that contains fool's gold without assuming that the world you entertain contains gold. As Kripke puts it, the allusion to gold fixes the reference of "fool's gold," but doesn't give the meaning of "fool's gold." (Mill self-consciously uses the same reference-fixing device to get at the possibilities of sensation that he identifies with physical things: see Mill, 1979, 197.)

¹¹ The severe difficulty of sharply defining the relevant kind of coherence is, I believe, what G.E. Moore is getting at when he says: "I have, no doubt, conclusive reasons for asserting that I am not how dreaming; I have conclusive evidence that I am awake: but that is a very different thing from being able to prove it. I could not tell you what all my evidence is; and I should require to do this at least, in order to give you a proof." (Moore, 1939, 149)

the experiences that other people have had, and that your dream experiences do not relate to other people's experiences. When I say that a given experience coheres with a totality of experiences to which it belongs, I mean that it relates to the other experiences in that totality in the ways your current experiences relate to the experiences that you and others have had (or similar ways), rather than the way that your dream-experiences relate to the experiences that you and others have had (or similar ways).

Normally, when an experience fails to cohere with the totality of experiences with which it co-exists, the non-cohering experience also fails to belong to a thinglike group of experiences. But it can happen that experiences belonging to a thinglike group fail to cohere with the totality of all experiences with which they co-exist.

Something like this happens when a hologram or other optical illusion makes it appear to multiple observers that there is a statue where there is in fact none. For a more extreme example, we can imagine a future where most people have neural implants with a "reality augmentation" function that induces highly realistic group hallucinations in select contexts. (A primitive precursor to this is already possible with so-called "smartglasses.") Suppose the implants are controlled by a central computer that induces realistic multi-modal experiences as of a unicorn in anyone who ventures into the vicinity of the seventh hole at Augusta National. When a group of people who have the implant go to that location, they all have perfectly coordinated experiences as of a unicorn on the green. Their experiences have the same intrinsic qualities, and relate to each other in the same ways, as they would if perceiving an actual unicorn.

Though their experiences cohere with one another, they don't cohere with the totality of *all* experiences that are apt to occur in their world, which include the experiences of people who approach the seventh hole without the neural implant, and the experiences they themselves (with the implant) would have if they tried to lead the unicorn off the course. Anyone who had a complete catalogue of all the experiences that are apt to occur in our world would notice that the unicorn experiences didn't fit into the totality in a coherent way. Despite their realism and mutual coordination, the unicorn experiences would stand out as aberrations, in a way that experiences of, e.g., an actual deer on the green would not. Spelling out the nature of the aberration in detail is what I've admitted I can't do, but that shouldn't discourage us from classifying the unicorn experiences as aberrant, any more than our inability to spell out necessary and sufficient conditions for dreaming as opposed to waking experience should discourage us from classifying certain experiences as dreams.

I said that a noumenalist should define physical things as noumenal powers to cause thinglike groups of experiences that cohere with the experiences that occur in the nearest ideal world, i.e. the ideal world that our experiences are most likely to occur in, conditional on their occurring in an ideal world. In the book, I motivate the need for the "nearest ideal world" clause by reference to the possibility of computer-generated worlds, like those described by Nick Bostrom in his discussion of "ancestor simulations," and the virtual world depicted in the *Matrix* movies.¹²

Suppose that such a computer actually exists in a laboratory at the University of Vermont, and that it's programmed to cause the sort of experiences that would be apt to

¹² For Bostrom's ancestor simulations, see Bostrom (2003).

occur in Tolkien's Middle Earth universe, if such a universe existed. The computer has the power to cause many thinglike groups of experiences that cohere with the other experiences that the computer has the power to cause. Among these are thinglike groups of experiences as of a city with the features Tolkien attributes to Minas Tirith. Are these experiences veridical perceptions of a real, physical city?

If noumenalists count as veridical any experience that coheres with some worldlike totality of experiences, they'll have to say that the Minas Tirith experiences are veridical. This is clearly the wrong thing to say, since Minas Tirith doesn't exist.

Should noumenalists require for veridicality that the Minas Tirith experiences cohere with *every* worldlike totality of experiences? Since the computer-generated experiences of Minas Tirith don't cohere with *our* experiences (yours, mine, the UVM computer scientists, etc.), this would spare noumenalists the embarrassment of saying that Minas Tirith actually exists. But it would also prevent them from saying that Burlington, Vermont exists, since our experiences of Burlington don't cohere with the worldlike totality whose constituent experiences the Middle Earth simulator has the power to cause.

What noumenalists should say is that the people in the Middle Earth simulator speak truly when they utter the words, "Minas Tirith is a real physical thing." It's just that by "a real physical thing," they mean something with the power to cause appropriate thinglike experiences that cohere with experiences in the ideal world that *their* experiences most likely to belong to, conditional on their experiences belonging to an ideal world. Since their experiences do not cohere with the experiences that make up the ideal world that *our* experiences most likely to belong to, we can also speak truly when we say, "Minas Tirith is not a real physical thing." We're just speaking different languages.

The oddball minds that DeRosset describes are much like the inhabitants of the Middle Earth simulator, or people who are born and live their whole lives with reality-augmenting brain implants of the type described earlier. The main difference is that the oddballs have no physical manifestation (being disembodied sentiences). It doesn't matter. The important point is that their experiences fail to cohere with the rest of the possible experiences in the ideal world that our own experiences most likely to belong to (in the relevant conditional sense). This is what gives noumenalists the right to deny that there are any real cantaloupes or unicorns answering to the oddballs' experiences. It's what lets noumenalists say that the oddballs are victims of mass hallucination.

What the noumenalists can say phenomenalists can say too. We just say it in terms of possibilities for experience, rather than in terms of experience-causing powers. Oddball minds are possible, including minds whose possibility entails that if certain streams of consciousness were to occur, they would include fully realistic and coordinated experiences as of unicorns and interstellar melons. If oddball minds actually exist (which we can't rule out), then the nearest ideal world includes all their unicornish and melonish experiences and the thinglike groups they constitute. None of that changes the fact that their unicornish and melonish experiences fail to cohere with all the other experiences that occur in the ideal world that our own experiences are most likely

to belong to. This is true, even if the oddballs' experiences somehow *do* cohere with the experiences that occur in the ideal world that *their* experiences are most likely to belong to (similar to the Middle Earthlings). That's why the possibility of oddball minds doesn't saddle the phenomenalist with unicorns.

3 Egocentrism

My response to the Demarcation Argument brings to the fore the crucial role of our own (or, more cautiously, *my* own) experience in the phenomenalist analysis of physical things, or any analysis in terms of experience-related modalities (counterfactuals, probabilities, causal powers, etc.). It's important to understand exactly what this role is, since a misunderstanding on this score can make phenomenism appear vulnerable to a very serious objection—one that does, in fact, knock the bottom out of Berkeleyan idealism (as well as contemporary panpsychism: more on this later).

According to Berkeley, physical things are literally made of experiences; an apple, for instance, is a combination of experiences occurring in the minds of anyone perceiving the apple (including God's mind). This commits Berkeley to saying that apples (or at least the ones that exist in our world) couldn't exist in a scenario where there were no minds or conscious experiences.

Berkeley was happy with this implication, because he erroneously thought he could prove that apples and all other physical things *are* mind-dependent in the way his theory implies. The truth is that we have every reason to think that apples could exist in a mindless world, and therefore every reason to reject Berkeley's metaphysics.

Despite the role that actual experience plays in the phenomenalist theory I favor, the theory doesn't have the bad Berkeleyan implication.

In my view, the physical facts reduce to facts about phenomenal probabilities, i.e. probabilities for certain purely experiential states of affairs to exist, conditional on the existence of other purely experiential states of affairs. That there actually is a physical world that actually has the specific physical features it does reduces to the fact that, given all the phenomenal probabilities that hold in our world, a specific logically possible ideal world is the one with the highest probability of existing and containing *my* experiences—the ones I actually have—conditional on there existing *some* ideal world that includes *my* experiences.

In the book, I tend to speak of “our” experiences rather than “my” experiences; this reflects my optimism that the ideal world most likely to exist conditional on there being an ideal world with all my actual experiences = the ideal world most likely to exist conditional on there being an ideal world with all *your* actual experiences.

Fumerton suggests that one way to understand what I'm doing in the book is announcing how I propose to use the phrase “the physical world,” and inviting others to use it the same way, in the defeasible but (I hope) reasonable expectation that if they accept my invitation, we'll all be using “the physical world” to refer to the same thing. (Or, more or less the same thing: we can admit a certain amount of variation

arising from marginal intersubjective differences in the concept of a physical thing: see pp. 137-38.)¹³

This is exactly how I want to be understood, and Fumerton's comparison with demonstrative and indexical reference is spot-on. It's like I'm inviting you to join me in using "human being" interchangeably with "the species I belong to," in the expectation that if you accept my invitation, we'll refer to the same thing by "human being" (allowing for a certain amount of variation that might arise from intersubjective differences in the concept of a species).

My theory assumes that there are logical possibilities for various ideal worlds that include all my actual experiences, and that one of these possibilities is more likely than the others to be realized, conditional on one of them being realized. If there is no logically possible ideal world satisfying this condition, then, in my view, there is no physical world, and my utterances of "the physical world" fail to refer.

The logical possibility of the relevant ideal world does not entail the actual existence of my experiences; it doesn't entail the actual existence of *any* experience. The relevant ideal world would have been a logical possibility, even if neither I nor my experiences had ever actually existed. Likewise, the existence of possible worlds where Julius Caesar lives to the age of 90 doesn't depend on Julius Caesar's having actually existed (that is, having existed in our, actual, world). It might have been difficult or impossible for us to refer to such worlds if Caesar had never actually existed, but the existence of a logical possibility doesn't depend on our being able to refer to it.

I mention my actual experiences in order to fix the reference of "the nearest ideal world," but that doesn't mean that "the nearest ideal world" refers to something that depends for its existence on my experiences actually existing. It refers to something that depends for its existence on my actual experiences being logically possible, but that is hardly objectionable, since it's logically necessary that they are logically possible.

Here is something that I don't say, but if I did would commit me to saying that apples are mind-dependent phenomena:

Apples exist $\rightarrow (\exists x)(\exists y)(x = \text{my experience} \ \& \ y = \text{the nearest logically possible ideal world containing } x \ \& \ \text{coherent thinglike experiences of apples occur in } y)$.

Rather, what my theory commits me to saying is:

$(\exists x)(\exists y)(x = \text{my experience} \ \& \ y = \text{the nearest logically possible ideal world containing } x \ \& \ (\text{apples exist} \rightarrow \text{coherent thinglike experiences of apples occur in } y))$.

¹³ Fumerton suggests that an additional source of variation might be differences in the types of experience by which different observers perceive physical things. In the book, I consider this possibility (pp. 139-43), and argue that it arises only in connection with hypothetical experience so radically unlike ours that, despite its orderliness, we can't relate it to ours in any intelligible way; depending on the details, beings who have such experience are best understood as not perceiving the physical world we perceive, or perhaps even as not perceiving a physical world at all. Less radical differences need not induce any difference in reference of "the physical world," since we can construe them on the model of beings whose visual experience is phenomenally color-inverted relative to our own; this is also how I would construe the experiences of Smithson's "subjects whose experiences are systematically stretched or compressed relative to ours."

This does not imply that apples depend for their existence on my having experience (or on my existing, or on any mind existing or having experience).

Though I don't discuss this in the book, it's interesting to consider scenarios where no logically possible ideal world satisfies the criteria whose satisfaction is, in my view, equivalent to the existence of the actual physical world.

One such scenario is where the probability of an ideal world that includes our actual experience is arbitrarily low. This is how it is, if all experience occurs purely at random.

We can also imagine scenarios where the probability of an ideal world that includes our experience is not arbitrarily low, but where such an ideal world can arise only at random, and multiple such worlds are equally likely (or equally unlikely) to arise. An example is a scenario where our experiences occur in a computer simulation run on a machine that doesn't have the power to generate enough experience to constitute an ideal world, perhaps due to constraints on computing power. (Maybe the computer only has enough power to generate experiences of things viewed from certain specific standpoints, as in what David Chalmers calls a "just-in-time" simulation.¹⁴) In this case, the likelihood of an ideal world that includes our experience might not be infinitesimal, but (to put it loosely) no worldlike completion of our experience is more probable than various alternative worldlike completions.

In a more intriguing case, one's use of "the physical world" fails to refer, because some of one's experiences are apt to belong to one worldlike totality, and others apt to belong to another. This is one way to understand Harry Potter's situation. Arguably, none of Harry's experiences cohere with the sum-total of all the experience that's apt to occur in his universe. Wizard experiences of Hogwarts prevent Muggle experiences of a certain Scottish valley from cohering with the totality of all possible experiences, and vice versa: the Wizard and Muggle experiences cast each other in the role of mass hallucinations. Still, Harry might choose to employ "the physical world" as an implicitly relational term, sometimes using it to refer to a subset of possible experiences that includes all Wizard experiences but excludes various Muggle experiences, and sometimes to a subset of possible experiences that includes all Muggle experiences but excludes various Wizard experiences.

Depending on unstated details about the experience-causing powers of the machinery that sustains the virtual world depicted in the *Matrix* movies, Neo and his friends can adopt the same linguistic policy, if they accept my theory. For Neo, there is no such thing as the physical world, but if he chooses, he may still speak of *a* physical world. For example, there is a totality of phenomenal probabilities that makes a certain ideal world the one that Neo's 21st-century-esque experiences are most likely to belong to, conditional on their belonging to a worldlike totality of experiences, and there is a distinct totality of phenomenal probabilities that makes a different ideal world the one that Neo's 28th-century-dystopian experiences are most likely to belong to, conditional on their belonging to a worldlike totality of experiences. Neo speaks truly when he says "there is a world where there are sushi bars" and also when he says "there is a world where there are no sushi bars." This is so, despite the fact that he can't truly say, "the world contains sushi bars" or "the world does not contain sushi

¹⁴ (Chalmers, 2010, 476).

bars” (at least, not without implicitly relativizing his utterances to relevant subsets of phenomenal probabilities).¹⁵

4 Time, Space, and the Causal Theory of Objects

One thing that distinguishes my theory from earlier phenomenalist theories is the use it makes of the concepts of time and space. This is not an idle novelty, but an innovation necessary to remedy a major shortcoming of traditional phenomenalist theories, namely their failure to account for the objectivity or intersubjective availability of physical things (a problem that traditional phenomenism inherited from the classic idealism from which it evolved).

Central to our concept of a physical thing is that it’s the sort of entity that more than one person can perceive. Causal theorists of the physical (such as Kant, and, as Fumerton points out, Hylas in the last stages of his resistance to Philonous) have no trouble accounting for this feature of physical things. According to causalists, for two people to perceive the same physical thing is just for one and the same thing to cause experiences in both people. (Actually, it’s not quite as simple as that, but some account along these lines is available to those who reduce physical things to possessors of suitable experience-causing powers.)

The ease with which the causal theory accounts for the objectivity of physical things is, I believe, its main advantage over phenomenism. Unless phenomenists can give an equally satisfactory account of physical objectivity, we should prefer a causal theory of objects to phenomenism.

Phenomenalists don’t deny that physical things (or many of them) have experience-causing powers, or that when Jack and Jill perceive the same thing, there is something that causes both of their experiences. But according to phenomenists, the existence of this common cause reduces to facts about possibilities for various experiences. In the phenomenist view, the common cause—whatever physical entity Jack and Jill perceive—just *is* a possibility for experiences to occur in certain ways, including the ways they occur in Jack and Jill. Ultimately, for Jack and Jill to perceive the same thing is for each of them to have experiences that belong to the same coherent thinglike group of experiences (in the nearest ideal world).

But when do different subject’s experiences belong the same thinglike group?

Berkeley thought that physical objects were thinglike groups of actually occurring experiences. According to him, when Jack and Jill see the same pail of water, what makes their experiences part of the same pail-of-water-like group of experiences is that they bear a certain resemblance to each other and relate to each other in a mutually dependent way, such that it’s not an accident that the experiences resemble each other, and not an accident that if Jill has a visual image of the pail getting overturned, Jack does too.

¹⁵ Neo might be tempted to say that only the dystopian world *really* exists, on the grounds that what happens in the 21st century world depends on what happens in the dystopian world, but not vice versa. If he yields to this temptation, he’ll find himself in a skeptical quandary, since for all he knows, what happens in the dystopian world asymmetrically depends on what happens in a world that relates to it as it relates to the 21st century world.

Mill essentially co-opts Berkeley's account, substituting possibilities for thinglike groups of experiences for actually occurring thinglike groups. For Jack and Jill to perceive the same pail of water is for their experiences to realize a propensity for experiences that (a) occur in patterns suggestive of a pail of water, and (b) stand to each other in suitable relations of counterfactual dependence.

This account is inadequate. Suppose Jack lives on Earth and Twin Jill lives on a Twin Earth that is qualitatively and dynamically indistinguishable from Earth. As Jack observes a pail of water on Earth, Twin Jill observes an indistinguishable pail of water on Twin Earth. Since the Earthly and Twin Earthly pails are indistinguishable, and since any changes that occur on Earth are paralleled by indistinguishable changes on Twin Earth, Jack's and Twin Jill's experiences stand in exactly the same relations of counterfactual and probabilistic interdependence as Jack's and (Earthly) Jill's. If, like Berkeley and Mill, we say that their experiences standing in those relations suffices for their being perceptions of the same pail, we have to say that Jack and Twin Jill perceive the same pail. But they don't.

Before turning to my solution to this problem, let me go over a couple of non-solutions. (I don't discuss these in the book, but sometimes it's easier to see where an author is coming from if you're shown the hidden traps that his approach is meant to avoid.)

One difference between the scenario where Jack and Jill see the same pail of water and the one where they see different but indistinguishable pails is that in the second, but not the first, there's a possibility for a coherent perception of two indistinguishable pails of water. Why not say that two experiences are perceptions of the same pail of water just in case they occur in a world where there's no possibility for a coherent third experience of two pails indistinguishable from the pail presented in each of the first two experiences?

Because this would imply that if Twin Earth exists as described, then when Earthly Jack and Earthly Jill see the same pail, their experiences aren't perceptions of the same thing, since in this scenario there is the possibility for a coherent third experience of two pails indistinguishable from the one that Earthly Jack and Earthly Jill perceive.

Well, then why not say that Jack and Jill perceive the same pail just in case there's a possibility for a coherent experience of both of them looking at the same pail? The possibility for this "bird's-eye view" experience exists, regardless of what's going on over on Twin Earth.

There are two problems with this proposal. First, it assumes that Jack and Jill have perceivable bodies; this is a problem, because the experiential possibilities that form the phenomenalist's reduction base for physical reality include the possible experiences of ideal observers, who lack perceivable bodies. (Ideal observers lack perceivable bodies to prevent them from interfering with the physical world they perceive: any world that contained enough perceptibly embodied observers to perceive all of its physical contents would differ from our world physically, due to the presence of all those bodies.) Second, the proposal *assumes* that Jack's experiences are experiences of the same pail as Jill's experiences, rather than, say, Jack's experiences being experiences of the Earthly pail, and Jill's being experiences of the Twin Earthly pail. If this assumption is justified, we need to seek a justification for it elsewhere.

My solution to the objectivity problem is to say that Jack and Jill see the same pail because their experiences occur at the same location in a common time and space. If Jack's and Twin Jack's experiences occupy different regions of the same spacetime, and Jack's and Jill's experiences occupy the same region of that spacetime, that is a basis for saying that Jack and Jill perceive the same pail of water, and that Jack and Twin Jack do not.

This might seem like a wild proposal, but it actually looks quite tame when viewed against the background of the modern scientific understanding of time and space. That said, looking back, I can see that my discussion of time and space (pp. 76–90, 185–93) is rather heavy going, and I regret dragging my readers deeper than necessary into the adjacent weeds. Let me take this opportunity to clarify my position on this, both in terms of its purpose and its overall shape.

In the now standard scientific way of understanding time and space—abstracting away from all nitty-gritty details—a thing's spatiotemporal location is given by the spacetime coordinates it gets assigned in an overall assignment of spacetime coordinates to events that has the property of making sense of the world—i.e., explaining the world's *prima facie* complexities in terms of a small number of intelligible natural laws. How we arrived at a system of coordinates for physical events that has this property is a fascinating story from the history of science whose details need not detain us here. (The weeds beckon; we heed them not.) The important thing to know is that you discover such a system by trying out different coordinate schemes, until you hit upon one that has the desired property of being such that when you describe the world in terms of that scheme, you describe it as a world thoroughly governed by simple, intelligible laws.

There is nothing in this method of assigning coordinates to events to prevent us from applying it to experiential events. This is true, even if experiential events are not physical events. As long as there's a scheme for assigning coordinates to conscious experiences that makes sense of experiential reality, we can say that experiences that receive the same coordinates in that scheme occur at the same location in experiential spacetime, and those that receive different coordinates occur at different locations in experiential spacetime.

In effect, you establish an experiential spacetime (or “ideal” spacetime, as I call it in the book) by construing actual and possible experiences as a phenomenal model of a physical world, and then assigning coordinates to the model's parts (various experiences) in a way that explains the model's phenomenal complexities in terms of simple, intelligible laws. (These laws, which I call “laws of experience,” basically tell experience to happen in a way that presents the appearance of a physical world obeying simple laws of physics.)

Here's an analogy. Suppose you use a computer to run a model of a galaxy. Various machine states of the computer's hardware model different parts of the galaxy. If we want, we can assign these states coordinates in the spacetime in which the contents of the computer lab and the world at large have coordinates (the world of tables, chairs, coffee mugs, the cars in the parking lot outside, etc.). But we can also assign the machine states coordinates based on the roles they play in the model they instantiate. For example, if one state models a star S_1 , and a second state models a star S_2 , and a third models a star S_3 that occurs twice as far from S_1 as S_2 , we can assign the machine

states coordinates in such a way that the Euclidean difference between the first state's coordinates and the third state's coordinates is twice the difference between the first state's coordinates and the second state's coordinates.

We can do all this regardless of whether there actually is a galaxy answering to the computer model. Even if there is no such galaxy, we can look at the model itself, and study it as a phenomenon in its own right. One way to make sense of the model's behavior is by assigning its constituent states the same coordinates that we would assign to various features of a galaxy answering to the model, if we were to assign coordinates to that galaxy's features in a way that made sense of the galaxy's behavior in terms of a small number of simple physical laws. In practice, it might never occur to us to coordinatize the computational states this way if we didn't think of the states as parts of a model of a galaxy, but such a coordinatization is one we *could* hit upon even if we had no notion of galaxies or physics.

We can think of the experiences that occur in the ideal observers of a physical world as a *phenomenal* model of a physical world. A computer model (of a galaxy, or whatever) is an input-output architecture that takes physical states (of a computer) as inputs and gives physical states as outputs; a phenomenal model (e.g., of a physical world) is an input-output architecture that takes phenomenal states as inputs and gives phenomenal states as outputs. Just as we can assign coordinates to computational states in a way that makes sense of the complexities of the computational model they belong to, we can assign coordinates to phenomenal states in a way that makes sense of the complexities of the phenomenal model they belong to. We can do this, regardless of whether there's any physical world answering to the model.

Ideal spacetime is the spacetime in which experiences have the coordinates that they get assigned in a sense-making coordinatization of a phenomenal model that comprises the experiences that occur in the most probable worldlike totality of experiences that includes our own experiences. Experiences with the same coordinates in that coordinatization have the same location in ideal spacetime; experiences with different coordinates in the coordinatization have different locations in ideal spacetime.

The possibility of a sense-making assignment of coordinates to experiences is guaranteed by the possibility of a sense-making assignment of coordinates to physical events. This is because physicists assign coordinates to physical events based on *perceived* (or at any rate perceptible) natural regularities. To get a coordinatization of experiences that makes as much sense of experiential reality as the physicists' coordinatization of physical events makes of physical reality, we need only substitute the relevant physical regularities with conscious perceptions of those regularities. If the physical coordinatization is possible, so too must be the phenomenal coordinatization.

Ideal spacetime is not physical spacetime, but the two are related. The locations of physical things in physical spacetime partly depend on the locations of experiences (or possible experiences) in ideal spacetime. Indeed, the very existence of physical things depends on the ideal spacetime locations of experiences in the nearest ideal world, since in my view, physical things are possibilities for groups of experiences whose constituents occur at the same location in ideal spacetime.

However, it doesn't follow that an experience's location in ideal spacetime is the same as its location (if any) in physical spacetime, any more than a fabric swatch's location in RGB color space is the same as its location in physical space. Nor does it

follow that experiences with the same location in ideal spacetime have the same location in physical spacetime. Assuming it makes sense to locate conscious experiences in physical spacetime, Jack's experiences of the pail presumably occur at a different location in physical spacetime from Jill's (Jack's somewhere in his head, Jill's somewhere in hers). This doesn't prevent their experiences from occurring at the same location in ideal spacetime. They can have the same location in ideal spacetime, despite having different locations (if any) in physical spacetime, just as identically-colored swatches can have the same location in RGB color space without having the same location in physical space. Or, to revert to the computational analogy, machine states of the galaxy-modeling computer might have different locations in the physical spacetime of the computer lab, despite having the same location in the spacetime of the model (e.g., the machine states might redundantly model the same star, or non-redundantly model different global features of the same star).

Zooming back out to the big picture: there is, I assume, a worldlike manifold of experiences that is the one our experiences have the highest likelihood of belonging to, conditional on their belonging to a worldlike manifold of experiences. These experiences have various phenomenal qualities, and stand in various relations, including relations of dependence (understood causally, counterfactually, or probabilistically). Since these qualities and relations include some that mirror the physical qualities and relations that physicists use to define a common spacetime for physical events, they allow phenomenalists to define a common (though distinct) spacetime for phenomenal events.

Fumerton asks how I understand the claim that there exists a physical sphere somewhere in the universe. I understand it to mean that in the nearest ideal world, there is a group of experiences that includes (among others) experiences of the sort one would typically have if (a) viewing the surface of a spherical object, (b) simultaneously viewing the surface and interior of a spherical object (X-ray vision style), (c) manually grasping a spherical object (if the object is firm enough to be grasped), (d) exploring the interior of a spherical object (if it's hollow, or non-solid, like a gas giant), or (e) walking around on the surface of a spherical object (if the object's physical nature permits that), where all the experiences in the group are connected by relations of counterfactual or probabilistic dependence, *and* they have identical or suitably overlapping locations in the ideal spacetime in which they occur.

If the sphere exists in a physical universe that lends itself to the sort of coordinatization by which physicists define physical spacetime, and in terms of which physicists describe an object as having a certain size and shape, and standing at certain spatiotemporal distances from other objects, then experiences of the sphere exist in a phenomenal universe (or rather possible universe) that lends itself to the same sort of coordinatization. If the sphere does not exist in a universe that lends itself to the relevant kind of coordinatization, then it's unclear in what sense it has a determinate size, shape, or location; in such a scenario, it's not even clear that it makes sense to speak of a spherical object, as opposed to a mere tendency for experiences to occur in ways that give the impression of such an object.

5 Metaphysical Realism

In his contribution, Smithson offers a characteristically trenchant summary of the realist challenge to phenomenalism. Smithson focuses on alleged explanatory shortcomings of phenomenalism in comparison to metaphysical realism, and these will also be my focus in what follows.

Broadly speaking, metaphysical realism combines the commonsense view that there are physical things with the claim that the physical facts of our world do not reduce to mental facts.¹⁶

The most common type of metaphysical realism today holds that the physical facts of our world don't reduce to *anything*. This is what we might call antireductionist metaphysical realism.¹⁷ A less common and more ambitious type of metaphysical realism holds that the physical facts of our world all reduce to more basic non-mental facts. The prime example of this is ontic structural realism, which proposes to reduce the physical facts to "structural" facts characterizable in purely logical and mathematical terms (augmented, in most versions, with some term for a relation of non-logical dependence, like causation or conditional probability).

Antireductionist realism poses no distinctive threat to phenomenalism. Antireductionist realists simply deny that the phenomenalist's reductive program, or any attempt to reduce the physical, can succeed. Maybe they're right, but to establish that they need to come up with some argument against reductionism about the physical. The bare assertion that physical phenomena are irreducible is not such an argument.

Still, antireductionist realism is so common these days, and reductionist theories of the physical so rare, that one can be forgiven for thinking that the former must have some distinctive *prima facie* advantage over the latter. Before considering structural realism, which (as I believe Smithson would agree) is the real threat to phenomenalism in this neighborhood, let me dispel this common misconception.

The main argument for phenomenalism is that it offers the most parsimonious defensible explanation of the regularity of experience, understood as the *prima facie* non-random and intersubjectively coordinated character of much actual experience. Briefly, the phenomenalist explanation is that what ultimately explains the fact that experience seems to occur in non-random and intersubjectively coordinated ways is that it really *does* occur in non-random and intersubjectively coordinated ways—a hypothesis that I couch in terms of objective probabilities for certain experiential states of affairs to exist, conditional on the existence of other experiential states of affairs.

Antireductionist realists offer a different explanation of the regularity of experience. It goes like this. We have physical bodies that are embedded in a physical environment. Interactions between these bodies and their environment cause the bodies (or minds associated with them) to have various conscious experiences. This is surprising, but given that it's true, it's not surprising that similar bodies would have similar

¹⁶ Some define metaphysical realism more broadly as the view that physical phenomena are real and do not in general depend for their existence on the existence of minds; on this definition, phenomenalism is a kind of metaphysical realism.

¹⁷ Everyone agrees that many physical facts reduce to simpler physical facts. The antireductionist claim is that not *all* physical facts reduce, and none reduce to non-physical facts.

experiences in similar environments, or that the same body would have similar experiences when affected by similar physical events. The regularity of experience is due to the propensity for similar bodies to have similar experiences when subject to similar physical influences. (As for what the bodies and environment *are*, antireductionist realists don't have much to say about that, since they think that physical phenomena are irreducible to anything more basic.)

This causal explanation of the regularity of experience is reasonable, and phenomenologists should accept it. Given what we know about how our bodies work and how their workings bear on our mental lives, *everyone* should accept it. But the cogency of the causal explanation does not preclude the possibility of an additional, deeper, and more ontologically parsimonious explanation, which is what the phenomenologist explanation aspires to be.

It would be silly to deny that when Jack and Jill perceive the same tree, we can explain the similarity and prima facie coordination of their experiences by pointing out that the tree has similar effects on their sensory surfaces, which have similar effects on their brains, which result in similar conscious experiences. A phenomenologist accepts all this, and adds a reductive analysis of the tree, sensory surfaces, and brains. According to phenomenologists, these are all possibilities of sensation.

Given this analysis, we have a further, and in a sense deeper, explanation of the regularities in Jack's and Jill's experiences of the tree. The prima facie coordination and non-randomness of the experiences is due to probabilities for there to be coherent tree-like experiences, conditional on there being all the other experiences that occur in Jack, Jill, and anyone else observing the tree. The objective probability that Jack's and Jill's mental lives include these tree experiences, conditional on their mental lives having all the other conscious features they have, is (much) greater than chance. Maybe it's surprising that Jack and Jill have any experience at all—the existence of consciousness is something of a mystery—but given that they do, it's not surprising that their experiences conform to the objective probabilities governing experience.

The phenomenologist explanation of the regularity of Jack's and Jill's experience is compatible with the causal explanation that antireductionist realists offer. It's also more parsimonious than the causal explanation. The antireductionist must grant that there are objective probabilities shaping Jack's and Jill's experiences, since the alternative is that their experiences occur randomly and independently of one another, which would make their regularity inexplicable. The phenomenologist just points out that we can get an explanation of the experiential regularities by appealing directly to these probabilities.

Phenomenologists don't have to stop there. The existence of the relevant probabilities also has an explanation: whatever explains why the tree has the features of it that Jack and Jill perceive, together with whatever explains why Jack and Jill have the physiological features they do. The tree's visible surface consists of plant cells with various physical properties. The existence of these cells with these properties reduces to the existence of suitable phenomenal probabilities. Likewise for the molecules that constitute the cells, the atoms that constitute the molecules, and (if we are going to be scientific realists) the subatomic particles that constitute the atoms. Similar remarks apply to the cells, molecules, atoms, etc. that constitute Jack's and Jill's bodies. If we reach a point where no further phenomenologist explanation is possible, that's because

we've reached physical phenomena that don't reduce to more basic physical phenomena. But these have no explanation according to antireductionist realists either.¹⁸

Smithson points out that “phenomenalists cannot, in general, explain the [regularity] of possible experiences” in terms of the experience-causing propensities of physical objects together (if necessary) with psychophysical laws that relate neural events to phenomenal events. That's correct: the general explanation appeals not to causes or psychophysical laws, but fundamental probabilities. In terms of causes and psychophysical laws, the only general-ish explanation a phenomenalist might give would be in terms of an infinite hierarchy of experience-causers, if such a hierarchy exists. But antireductionist realists can't give a general explanation for the regularity of *physical events* in *any* terms, unless it's a general-ish explanation that posits an infinite hierarchy of physical events. Antireductionist realists aren't racking up any points against phenomenalists, here.

Smithson also suggests that phenomenalist explanations of the regularity of experience are more complex than causal explanations, and this might be true, if “more complex” means harder to state in an explicit and compendious way. But this is no objection to the phenomenalist explanation. If phenomenalists have to use more words than antireductionist realists, that's only because they have more to say.¹⁹

Suppose someone gives a folk-psychological explanation of various patterns of human behavior. Now a reductionist *physicalist* comes along and proposes an analysis of psychological phenomena in terms of complex states of the human nervous system. The analysis is compatible with the folk-psychological explanation, and yields an additional—and deeper—explanation of the relevant behavior (in terms of the neural states that the folk-psychological states ostensibly reduce to).

Maybe physicalism is false, but it's no objection to it that *if* it's true, it provides an analysis of folk-psychological concepts in terms of complex neural states, and an explanation of human behavior in those terms. Likewise, it's no objection to phenomenalism that if it's true, it provides an analysis of physical concepts in terms of possibilities for complex experiential states, and an explanation of the regularity of experience in those terms.

The real threat to phenomenalism comes from ontic structural realism (“structuralism,” for short). Like antireductionist realists, structuralists give a causal explanation for the regularity of experience. But unlike the antireductionists, structuralists also

¹⁸ Dualist phenomenalists will have to complete their explanation of the regularity of Jack's and Jill's experience with fundamental psychophysical laws, but the same goes for all dualist explanations, including dualist realist explanations. (Physicalist phenomenalists have no need for fundamental psychophysical laws.) Smithson suggests that dualist antireductionist realists “put...psychophysical laws to more efficient use” than dualist phenomenalists. I disagree: they put them to the same use; only, phenomenalists also provide an account of the physical side of the psychophysical equation.

¹⁹ This is also my reply to Hedda Hassel Mørch's observation that while “[a]ccording to other theories, the physical world has...the same simple structure as posited by physics...[a]ccording to phenomenalism, the physical world consists in possibilities for a vast, perhaps even infinite, number of experiences...” Phenomenalists agree that the physical world has the simple structure posited by physics; it's just that we reduce this structure's elements to complex constructions out of possible experiences.

give an analysis of the physical. Structuralism is therefore in direct competition with phenomenism.²⁰

According to the most defensible versions of structuralism, physical facts reduce to facts expressible in terms of powers, dispositions, or probabilities. In this respect, structuralism resembles noumenalist and phenomenalist metaphysics. The difference is that where noumenalists reduce physical facts to facts about powers or dispositions to bring about experience, and phenomenalists to probabilities related to experience, structuralists reduce physical facts to facts about powers or dispositions to bring about, or probabilities related to, states of affairs characterized purely in logical, mathematical, and modal terms. Where phenomenalists see the physical world as an input-output architecture that takes possible experiences as inputs and gives possible experiences as outputs, structuralists see it as an input-output architecture that takes bits of input-output architecture as inputs and gives bits of input-output architecture as outputs.

Structuralism's ontic austerity is at once its greatest asset and its greatest liability. It's an asset inasmuch as it makes structuralism the most parsimonious analysis of the physical that has any hope of succeeding. *Every* analysis of the physical needs to posit logical and mathematical properties and relations, and some relation of non-logical dependence. Since structuralism posits *only* these, no defensible theory can be more parsimonious than it. At the same time, it's hard to get one's head around the idea that one can tell the whole truth about an apple in a language as meagre as the structuralist claims to be adequate for this purpose. Given that mathematics more or less reduces to set theory, such a description would basically say (1) that there are a certain number of things, (2) that these things belong to various sets, ordered series, and set-theoretic constructions thereof, and, (3) that the existence of various of these things depends in some non-logical way on the existence of various others.

Hard to get our heads around or not, there's no denying that, other things being equal, structuralism's austerity gives it an advantage over alternative analyses of matter, including phenomenism. One point to structuralism.

Another *prima facie* advantage of structuralism is its handling of nomicly imperceptible physical phenomena. In the structuralist scheme of things, imperceptible phenomena are just further structural features of the world (terms of various equations, as it were). Like all physical things, they can exist regardless of any associated experience-related powers or probabilities. (Structuralists do have to grant that macroscopic physical things have associated experience-causing powers, since that's how they explain the regularity of experience.)

In the book (pp. 59-62), I argue that it's a mistake to think that we can divorce physical things from experience-related powers or possibilities. In deciding whether to interpret a *prima facie* reference to an ostensible physical phenomenon as (1) a genuine reference to a real physical phenomenon, or (2) a mere manner of speaking (as in scientific antirealist accounts of things like quarks), the only compelling reason we can have for taking the first option is that the ostensible physical phenomenon

²⁰ By "structuralism," I mean specifically structuralist reductionism about the physical. This is the kind of structuralism we find in Russell (1927), Ladyman and Ross (2007), Sider (2011), French (2014), and Tegmark (2014). A more ambitious structuralism tries to reduce everything, including mental phenomena, to structure. (In her contribution, Hedda Hassel Mørch uses "ontic structural realism" (or "OSR") to refer to this stronger view.)

is something that is perceivable in principle (even if not in practice or perhaps even consistently with natural law). There are other reasons one might have for favoring (1) over (2), such as the much-discussed “no-miracles” argument, but these aren’t compelling, since there are *prima facie* countervailing reasons for favoring (2) over (1), such as the much-discussed pessimistic meta-induction.²¹

In my view, a quark is basically a schedule of probabilities for the occurrence of certain logically possible but nomically impossible experiences, conditional on the occurrence of other logically possible experiences. If, as I maintain, the belief that quarks are in-principle perceptible is the only thing that can compel us to believe there are quarks (rather than just a useful kind of quark-talk, like talk about “the average man” or “Coriolis forces”), then my account of quarks is defensible as long as we have compelling reasons to believe that there are quarks. (If we *don’t* have compelling reasons to believe there are quarks, there’s no pressing need to account for them.) More generally, any reason to doubt that phenomenalism can account for some alleged imperceptible thing is a reason to doubt (though not necessarily to deny) that there’s anything to be accounted for in the first place.

Nevertheless, some people may still prefer the structuralist account of quarks and such to the phenomenalist account. If the structuralist account really is preferable (despite my arguments to the contrary), then structuralism is superior to phenomenalism, other things being equal.

The problem, for structuralists, is that other things are not equal.

The great advantage of phenomenalism over structuralism is that it does justice to our pre-theoretical intuitions about what considerations are ultimately relevant to what physically exists.

Take the structure that structuralists identify with the physical universe we inhabit; call it S . If our world had been very different from how it actually is, physically (maybe different laws of physics, or different boundary conditions), it would, by structuralist lights, have been a different structure, call it S' .

We can imagine a world, W_S , in which structure S exists, but in which experience tends to occur in ways that are suggestive of a physical world very unlike ours: ways that you’d expect experiences to occur in a world with very different laws of physics from ours, or very different boundary conditions. For instance, we can imagine that in W_S , the only experiences apt to occur are coherent, thinglike experiences as of a world consisting of nothing but a handful of stars and some interstellar dust. We can also imagine a world, $W_{S'}$, in which structure S' exists, and in which experience tends to occur exactly as it tends to occur in our world: whatever experience-causing powers exist in our world exist in $W_{S'}$, and whatever experience-related counterfactuals or conditional probabilities hold in our world hold in $W_{S'}$.

²¹ See (Putnam, 1975, 73) and Laudan (1981). (For the reasons given in Magnus and Callender (2004), it seems to me that neither of these arguments succeeds.) Smithson imagines star-writing of putatively divine origin that announces the existence of in-principle imperceptible physical phenomena. I’m not sure how to fill in the details of this story so that it’s more reasonable to take the stars as evidence of in-principle imperceptible physical phenomena than to take the arguments against in-principle imperceptible physical phenomena as evidence for the unreliability of the stars. In any case, the mere possibility of such a scenario is no argument against phenomenalism; if it were, then similar scenarios would weigh against all theories.

Structuralism implies that W_S is exactly like our world, physically, and that $W_{S'}$ is radically unlike our world, physically. It seems to me that these are unacceptable implications.

If we want to discover the physical structure of our world, we have to start by figuring out what physical world we inhabit. How do we do that? Not by *a priori* reflection. We do it by having and thinking about experiences. We have experiences, we compare notes with other experiencers, we do some thinking about all this experience, and on that basis arrive at conclusions about our world's physical contents. At this point, we may hypothesize that the physical world has a certain structure to it—a structure that perhaps explains the lawlike behavior of its contents. But the hypothesis is only as plausible as the claim that there is a physical world with certain features, and that claim depends on facts about what experiences occur and are apt to occur.

In the book, I try to make the point with what I call the Shifting Basis Argument (pp. 68-71). I ask you to imagine that our world has some underlying nature that possesses experience-causing powers or in some other way serves as truth-maker of experience-related counterfactuals or conditional probabilities, but that this nature is constantly changing. The underlying nature shifts from moment to moment, from a Berkeleyan God to a population of Leibnizian monads to a Matrix-style supercomputer and so on. We can also imagine the underlying nature shifting from one structuralist structure to another: from S to S' to S'' . . . etc.

If you were somehow to learn that this was happening, would you conclude that no physical thing exists for more than a moment? That the shoes you put on this morning were different from the ones you removed at the end of the day? I believe you would not. If I'm right about that, you should agree that it's our world's experiential powers or propensities that decide its physical contents, not whatever abstract or quasi-abstract structure might underlie those powers or propensities.²²

Phenomenalism respects our intuition that when it comes to forming an opinion about the world's physical contents, experiential facts supersede structural facts. When we hold the structural facts fixed and vary the phenomenal probabilities, we change the physical facts, but when we hold the phenomenal probabilities fixed and vary the structural facts, we do not. It seems to me that this gives phenomenalism a decisive advantage over structuralism.

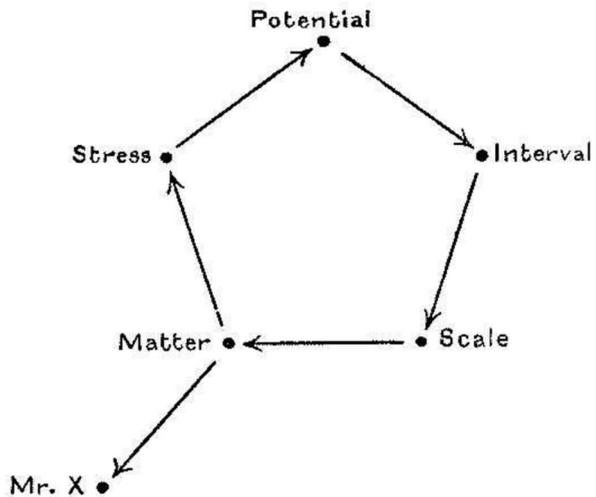
6 Panpsychism

Hedda Hassel Mørch has emerged as one of panpsychism's most clear-eyed and forceful proponents, and I welcome the opportunity her contribution affords to put phenomenalism into dialogue with panpsychism.

Panpsychism and phenomenalism are both mind-first theories of the nature of the physical, but the resemblance between the two ends there. According to phenomenologists, physical phenomena are fundamentally modal, consisting of possibilities for

²² I consider this to be the central intuition behind all broadly idealist metaphysics, as articulated by Leibniz (2007, 257), Berkeley (1982, §20), Mill (1979, 183-84), Ayer (1947, 178), Smithson (2017), and Adams (2021, 118-32).

Fig. 1 Cyclic definitions of physics + Mr. X



experience (understood as a certain species of probabilities, in my version of the theory); according to panpsychists, physical phenomena are fundamentally categorical, consisting of elementary “micro” experiences and constructions out of them. According to phenomenologists, the phenomenal character of actual and possible experience is highly relevant to the world’s physical contents; according to panpsychists, phenomenal character plays little if any role in determining the world’s physical contents (who knows what micro-experiences are like?). Last but not least, according to phenomenologists, there’s no contradiction in the claim that actual physical things could have existed even if there had never been any conscious experience; according to panpsychism, it’s metaphysically impossible for any actual physical thing to exist in a world devoid of consciousness.

One way to think of panpsychism and phenomenism is as alternative solutions to Arthur Eddington’s “cyclicity” puzzle. (A third solution is structural realism: more about this in a moment.) The puzzle arises from the prima facie circularity of physicists’ definitions of their basic vocabulary. Physicists define their basic vocabulary operationally, in terms of various measurement protocols: essentially, uses of various instruments to get certain outputs (“pointer readings,” in Eddington’s jargon) for certain inputs (“observations” or “experiments”). The problem is that the instruments and their inputs and outputs are themselves typically understood as physical phenomena possessing the very physical features that the operational definitions are meant to define (see Fig. 1, taken from Eddington’s *The Nature of the Physical World*²³).

One solution to this puzzle, which Eddington does not consider, is that the circularity of the physicists’ definitions is benign. Not all circular systems of definitions are circular in a bad way: some contain the resources to extract perfectly good non-

²³ (Eddington, 1929, 263).

circular definitions of the relevant vocabulary. The extraction method is the so-called Ramsey-Lewis method of defining theoretical terms.²⁴

Going into the details of a Ramsey-Lewis style solution to Eddington's puzzle would take us too far afield; suffice it to say that the result of applying the Ramsey-Lewis method to the *prima facie* circular definitions of basic scientific vocabulary yields definitions of elementary physical phenomena in purely structural terms.²⁵

Setting aside structuralism, there are two general approaches to solving Eddington's puzzle. Both start by challenging the assumption that we need to think of the instruments by reference to which we operationally define basic physical vocabulary in purely physical terms. What if we think of ourselves as scientific instruments? And what if we think of our "pointer readings" as *conscious* states, rather than (merely) physical states?

If we characterize Mr. X—a scientist conducting some experiment or making some observation—in purely physical terms, he's just like a weighing scale, in which case appending him to the chain of definitions doesn't break us out of the circle. But if Mr. X's "pointer readings" are non-physical states of Mr. X, then defining physics vocabulary in terms of Mr. X's pointer readings might break us out of the circle after all.

At this juncture, it would have been very natural for Eddington to have proposed a broadly Kantian solution to his puzzle. Instead of defining physical phenomena as entities that are apt to have certain physical effects on physical instruments, define them as entities that are apt to have certain *experiential* effects on conscious *minds*. Having taken this step, Eddington might have followed Mill in taking the further step of defining physical phenomena as tendencies for the world to give certain experiential outputs for certain experiential inputs (in effect, permanent possibilities of sensation).

But this is not the direction Eddington takes. Instead, the idea of treating Mr. X as a scientific instrument with experiential pointer readings inspires him to conjecture that various physical states of Mr. X *are* conscious experiences. The reasoning seems to be (1) human pointer-readings = relevant physical states of human brains, (2) human pointer-readings = various conscious states, so, (3) relevant physical states = conscious states. Since there's nothing about the physical states that occur in human brains that makes them seem more likely than other physical phenomena to be conscious,

²⁴ See Ramsey (1978), Carnap (1956), Lewis (1970), and Lewis (1972). The Ramsey-Lewis method takes a body of sentences that use the terms to be defined, conjoins those sentences, replaces all occurrences of the terms in the resulting conjunction with second-order variables of quantification (assuming the terms we're defining refer to properties or relations), and then generates so-called Ramsey sentences of the form, " $\forall\alpha(\alpha$ has property F [e.g., mass] iff there are properties and/or relations ϕ, ψ, χ etc. such that for all x, y, z etc., ... & $F=\phi \ \& \ F\alpha$)"

—where the ellipsis gets filled by an open-sentence that includes variables bound by the quantifiers plus other terms whose meanings are already known. Even if the terms to be defined occur in circular references to each other in the sentences we start out with, the Ramsey sentences we derive from the original sentences need not constitute a circular system of definitions.

²⁵ This is because the scaffolding of the relevant Ramsey sentences—basically, the sentences' non-quantificational vocabulary—consists solely of mathematical terminology and a primitive predicate for causation (or some other relation of non-logical dependence). A common complaint about structuralist metaphysics is that it's unclear that this scaffolding provides enough material to give the defined terms their intended meanings.

Eddington concludes that *all* physical phenomena are conscious: that (as he puts it) “the stuff of the world is mind-stuff.”²⁶

How does this solve the original puzzle? Not by replacing the standard operational definitions (in terms of physical inputs and outputs) with Kantian or Millian definitions (in terms of phenomenal inputs and outputs). Rather, the idea is to abandon the whole idea of identifying physical phenomena with the satisfiers of suitable operational criteria, identifying them instead with various parcels or configurations of mind-stuff. Once we divorce the project of defining scientific terms from that of discovering the intrinsic nature of physical reality, it doesn’t much matter whether the scientific definitions are circular, as long as they’re useful (which, as Eddington emphasizes, they certainly are).

From a Kantian or Millian standpoint, Eddington’s solution to the puzzle of cyclic definition looks rather clumsy. But the view Eddington arrives at—a kind of panpsychism—doesn’t stand or fall with its utility as a solution to Eddington’s puzzle, and many contemporary metaphysicians think it has independent merits that make it a consideration-worthy and even plausible account of the ultimate nature of physical reality.

Let me pause to state what I take to be the panpsychist theory most widely held today. This theory, also known as Russellian monism, consists of two claims: (1) fundamental physical phenomena are purely conscious phenomena (minds or experiences), and, (2) our own conscious experiences are combinations of the conscious experiences of the fundamental physical phenomena that make up our brains.

Insofar as the goal of panpsychism is to give an account of the nature of the physical, all it needs is (1). If fundamental physical phenomena are purely conscious phenomena, then the stuff of the physical world is indeed mind-stuff. The most basic physical phenomena (e.g., quarks, maybe) are elementary experiential phenomena; other physical phenomena (molecules, rocks, chairs, human brains, etc.) are complexes of elementary experiential phenomena.

The purpose of (2) is to solve the classic problem of explaining why conscious experience accompanies the physical processing that takes place in our brains. The panpsychist’s idea here is that if the elementary physical constituents of our brains just *are* conscious experiences of some sort (whatever experiences the panpsychists identify with quarks or whatever), then it’s not too much of a stretch to think that when those constituents combine together as they do in our brains, they collectively constitute the experiences that we introspect.

A main objection to panpsychism is that this *is* too much of a stretch. This is the so-called “combination problem” for panpsychism. I’m not going to go into it here, other than to register my pessimism for the prospects of a satisfactory solution. Let’s assume, for the sake of argument, that if the panpsychist identification of fundamental physical phenomena with conscious phenomena is correct, then we can identify our own conscious states with physical brain states.

Assuming this, Hedda Hassel Mørch is right to point out that panpsychism has an important advantage over phenomenalism. If our conscious states are just brain states, then there is no problem with saying that our conscious states have physical effects. On

²⁶ (Eddington, 1929, 276).

the other hand, if our conscious states are non-physical states, then we can't say that they have physical effects without either taking a bet at long odds against the causal closure of the physical world, or hypothesizing an implausible systematic overdetermination of certain types of bodily events by physical and non-physical causes. The standard view, which I endorse, is that if our conscious states are non-physical, we should *deny* that they have physical effects, explaining away appearances to the contrary with an epiphenomenalist account, by which our conscious states are inert by-products of physical states that cause both them and associated bodily behavior.

For most of the past sixty or seventy years, it was thought that the only plausible way to avoid epiphenomenalism was by *reducing* conscious states to physical states.²⁷ Such a reduction, if successful, would leave us free to say that conscious states have physical effects. Panpsychism opens up the prospect of avoiding epiphenomenalism *without* reducing conscious states to physical states. For panpsychism is not a reductionist theory of consciousness. It reduces complex conscious states (such as those that we introspect) to simpler conscious states; but, after all analysis is done, we're left with simple conscious states that don't reduce to anything. In particular, they don't reduce to physical states. Rather, they just *are* physical states. They are fundamental phenomenal states that are also fundamental physical states.

Unlike panpsychists, phenomenologists can't plausibly avoid epiphenomenalism without embracing a reductionist physicalist account of consciousness. For my part, I do not find this particularly disquieting, since epiphenomenalism doesn't strike me as particularly implausible, and the arguments against it strike me as weak.²⁸ That said, a lot of people find epiphenomenalism highly implausible, and many of these people also have grave doubts about reductionist physicalism. For them, panpsychism's ability to deliver a kind of mind-body monism free from epiphenomenalist implications is a major point in its favor, giving it an important advantage over phenomenology, which cannot deliver the same package of goods. One point to panpsychism.

As always, we need to weigh the advantages against the disadvantages. Two disadvantages stand out as particularly serious.

The first disadvantage of panpsychism is the same as that which I urged above in connection with structuralism. Suppose our world's fundamental categorical constituents are various microminds or microexperiences, as panpsychists contend. We can imagine a world where exactly these minds or experiences exist, related exactly as they are in our world, but in which they have experience-causing powers, or ground experience-related probabilities, just like those one would expect to characterize a world physically very unlike ours. Intuitively, such a world *is* physically unlike ours, but panpsychists have to say that such a world is physically indistinguishable from our world.

We can also imagine a world *W* characterized by the same experience-causing powers or experience-related probabilities that characterize our world, but where those powers or probabilities are grounded in microminds or microexperiences very different from those that ground them in our world (if panpsychism is true). Intuitively, *W* is physically indistinguishable from our world; if we somehow came to believe that

²⁷ For an illuminating explanation of this sociological fact, see Papineau (2001).

²⁸ See Pelczar (2021, 437–40).

our world's fundamental categorical constituents were the microminds or microexperiences that exist in *W*, we wouldn't conclude that our world's physical contents were different from what we previously thought. But this is the conclusion that panpsychists would be compelled to draw.

In short, when we hold the (alleged) microphenomenal facts fixed and vary the phenomenal probabilities, we change the physical facts, but when we hold the phenomenal probabilities fixed and vary the (alleged) microphenomenal facts, we do not.

Panpsychists can try to resist this by arguing that powers and probabilities require a categorical basis, and identifying the categorical basis of our world's experience-causing powers or experiential possibilities with something mental. This strategy has two shortcomings.

First, as I argue in the book (pp. 64–72), there is no good reason to think that powers and probabilities require a categorical basis; rather, there is good reason to think that they do not. For example, it seems clear that there could be a world indistinguishable from ours in terms of its categorical features, but in which radioactive elements have different half-lives from those they actually have. Such a world might be highly improbable, but it's not impossible. It's just a world in which a much larger than expected number radioactive particles "beat the odds." But if facts about probability supervened on categorical facts, such a world would be impossible.

Second, even if experiential powers or probabilities did require a categorical basis, it wouldn't matter to the world's physical contents what specific categorical basis this was (that was the upshot of the Shifting Basis Argument). As a corollary to this, it wouldn't matter whether the categorical basis was a congeries of panpsychist micro-experiences or something else. If panpsychist microexperiences exist, they might explain why our world has the experiential powers or probabilities it does, but in that case the right conclusion to draw is that the microexperiences explain why there is a physical world with certain physical features, not that they *are* that world.

Panpsychism's second major disadvantage to phenomenalism relates to its mind-dependent account of physical things.

A main advantage of phenomenalism over classic Berkeleyan idealism is that unlike Berkeleyan idealism, phenomenalism does not commit its proponents to a mind-dependent picture of physical reality. (Noumenalism enjoys the same advantage.) In the Berkeleyan scheme of things, it's impossible for the Moon to exist in a world devoid of consciousness; not so in the phenomenalist scheme of things, where the Moon's existence reduces to possibilities of sensation that can exist without being realized in the form of any actual experience.

When it comes to mind-dependence, panpsychists are in the same boat as Berkeleyans. Like Berkeleyan idealists, panpsychists have to regard the thought of the Moon existing in a world devoid of consciousness as a contradiction in terms. How could it fail to be, if, as panpsychists maintain, the Moon consists of fundamental physical constituents, and each of these constituents is a conscious mind or experience?

Let's put the challenge to panpsychism in the form of an argument:²⁹

- (1) We can imagine a world that is devoid of consciousness, but in which various actual physical things, such as the Moon, exist.

²⁹ The following discussion draws on Pelczar (2022).

- (2) What we can imagine is metaphysically possible.
- (3) So, it's metaphysically possible for the Moon to exist in a world devoid of consciousness.

There are ways to try to resist premise (2) of this argument, but panpsychists have good reason to go along with it. That's because a main motivation for panpsychism is to provide an alternative to reductionist theories of mind that try to analyze consciousness in physical terms, and a main reason to desire an alternative to this kind of physicalism is its vulnerability to the zombie argument, which relies on (2), or some more defensible refinement of it.³⁰

Instead, panpsychists object to (1). They say we don't know we can imagine a world physically indistinguishable from ours but devoid of consciousness. We know we can imagine a world that has all of the actual world's *structural* physical features but no consciousness, but (they say) we don't know that we can imagine a world that has *all* of the actual world's physical features but no consciousness, since (they say) for all we know, the actual world's *non-structural* physical features include intrinsic experiential features of fundamental physical things.³¹

The problem with this reply is that we *do* know we can imagine a world that has all the actual world's physical features but no conscious experience, since (a) we know we can imagine a world that has the same experience-causing powers or phenomenal probabilities as our world, but no conscious experience, and, (b) to imagine such a world is to imagine a world physically identical to ours but devoid of experience.

Panpsychists can't plausibly deny that we can imagine all our world's experience-causing powers or phenomenal probabilities existing in the absence of any actual experience: we need only imagine that all the powers or probabilities exist, but that none of them gets exercised or realized. (If panpsychists deny this, they have to explain why it's so easy to imagine powers or possibilities existing in scenarios where nothing realizes them.) If panpsychists want to resist the idea that the Moon could exist in an experienceless world, they have to argue that imagining a world that has all our world's experiential powers or possibilities is insufficient for imagining a world in which the Moon exists. In other words, they have to show that noumenalist and phenomenalist accounts of the physical are inadequate.

Can they show this? I hope not. In any event, the ball is in their court, and until they hit it back, I believe phenomenism maintains a comfortable lead.

Hedda Hassel Mørch distinguishes "mere mind-dependence" from "observer-dependence," and suggests that the former, which is all that panpsychism is committed to, isn't so bad. To say that physical things are observer-dependent is to say that they depend for their existence on being observed by some sentient observer; to say that they are mind-dependent is to say that they depend for their existence on the existence of *some* conscious experience, but not on experiences occurring in sentient observers of them.

³⁰ For possible refinements, see Chalmers (2002).

³¹ See Stoljar (2001) (though Stoljar doesn't advocate panpsychism) and (Chalmers, 2017, 28-29). In Stoljar's terms, the panpsychists' objection to premise (1) is that it's true only on a "theory-based" rather than an "object-based" understanding of the physical.

I agree that panpsychists aren't committed to an observer-dependent account of physical things, but I disagree that commitment to a "merely" mind-dependent account is acceptable. That rocks can exist in the absence of conscious experience seems to me as obvious as that they can exist in the absence of conscious observers of rocks. Maybe rocks or their subatomic constituents really are conscious, but if so, this is a contingent fact about them. A zombie rock is still a rock.

Author Contributions There is only one author.

Funding No funding was received for this study.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing Interests The authors declare no competing interests.

References

- Adams, R. M. (2021). *What Is, and What Is In Itself: A Systematic Ontology*. Oxford: Oxford University Press.
- Ayer, A. (1946-1947). Phenomenalism. *Proceedings of the Aristotelian Society*, 47, 163–196.
- Berkeley, G. (1710/1982). *A Treatise Concerning the Principles of Human Knowledge*. Indianapolis: Hackett Publishing. In K. Winkler (ed.).
- Bostrom, N. (2003). Are we living in a computer simulation? *The Philosophical Quarterly*, 53(211), 243–255.
- Carnap, R. (1956). The methodological character of theoretical concepts. In H. Feigl & M. Scriven (Eds.), *The Foundations of Science and the Concepts of Psychology and Psychoanalysis* (pp. 38–76). Minneapolis: University of Minnesota Press.
- Chalmers, D. (2010). The matrix as metaphysics. *The Character of Consciousness* (pp. 455–494). Oxford: Oxford University Press.
- Chalmers, D. (2002). Does conceivability entail possibility? In T. S. Gendler & J. Hawthorne (Eds.), *Conceivability and Possibility* (pp. 145–200). Oxford: Clarendon Press.
- Chalmers, D. (2017). Panpsychism and panprotopsychism. In G. Brüntrup & L. Jaskolla (Eds.), *Panpsychism: Contemporary perspectives* (pp. 19–47). Oxford: Oxford University Press.
- Eddington, A. (1929). *The Nature of the Physical World*. New York: Macmillan.
- French, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.
- Ladyman, J., & Ross, D. (2007). Ontic structural realism and the philosophy of physics. In J. Ladyman & D. Ross (Eds.), *Every Thing Must Go: Metaphysics Naturalized* (pp. 130–189). New York: Oxford University Press.
- Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, 48(1), 19–49.
- Leibniz, G. (1712/2007). Leibniz to Des Bosses, 16 June 1712. In B. C. Look, & D. Rutherford (Eds.), *The Leibniz-Des Bosses Correspondence* (pp. 255–257). New Haven: Yale University Press.
- Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13), 427–446.
- Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, 50(3), 249–258.
- Magnus, P., & Callender, C. (2004). Realist ennui and the base rate fallacy. *Philosophy of Science*, 71(3), 320–338.
- Mill, J. S. (1865/1979). *An Examination of Sir William Hamilton's Philosophy, and of the Principal Philosophical Questions Discussed in His Writings*. Toronto: University of Toronto Press.
- Moore, G. (1939). Proof of an external world. *Proceedings of the British Academy*, 25, 273–300.

- Papineau, D. (2001). The rise of physicalism. In C. Gillett & B. Loewer (Eds.), *Physicalism and Its Discontents* (pp. 3–36). New York: Cambridge University Press.
- Pelczar, M. (2021). Modal arguments against materialism. *Noûs*, 55(2), 426–444.
- Pelczar, M. (2022). The case for panpsychism: A critical assessment. *Synthese*, 400(4), 1–22.
- Putnam, H. (1975). *Mathematics, Matter and Method*. Cambridge: Cambridge University Press.
- Ramsey, F. (1929/1978). Theories. In D. Mellor (Ed.), *The Foundations of Mathematics and Other Logical Essays* (pp. 101–125). London and Henley: Routledge & Kegan Paul.
- Russell, B. (1927). *The Analysis of Matter*. London: Kegan Paul, Trench, Trubner & Co.
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Clarendon Press.
- Smithson, R. (2017). A new epistemic argument for idealism. In T. Goldschmidt & K. L. Pearce (Eds.), *Idealism: New Essays in Metaphysics* (pp. 17–33). Oxford: Oxford University Press.
- Stoljar, D. (2001). Two conceptions of the physical. *Philosophy and Phenomenological Research*, 62(2), 253–81.
- Tegmark, M. (2014). *Our Mathematical Universe: My Quest for the Ultimate Nature of Reality*. New York: Alfred A. Knopf.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.