

DEFENDING PHENOMENALISM

BY MICHAEL PELCZAR*

According to phenomenalism, physical things are a certain kind of possibility for experience. This paper clarifies the phenomenalist position and addresses some main objections to it, with the aim of showing that phenomenalism is a live option that merits a place alongside dualism and materialism in contemporary metaphysical debate.

Keywords: metaphysics, idealism, phenomenalism.

I. THE MILLIAN PICTURE

We all have perceptual experiences, which, taken together, present a subjective appearance of objects and events existing in a common time and space. In Leibniz's famous image, our experiences are like different perspective-drawings of the same landscape. They are, John Foster puts it, world-suggestive.¹

Ordinarily, we attribute the world-suggestiveness of our experiences to the fact that we all inhabit the same world, encounter objects in a common space, and witness events in a common time.

J.S. Mill thought that this way of thinking, while correct as far as it goes, misses out on a deeper truth. Yes, we have bodies with such-and-such physical features, embedded in such-and-such physical environments, and, yes, there's an explanation for the regularities in our experience to be found in all that. But, at a more basic level, the world we perceive doesn't *explain* the world-suggestive quality of our experiences: it *is* the world-suggestive quality of our experiences, or rather: it's the tendency for experiences to occur in a world-suggestive way, given that they occur at all. In Mill's view, physical things are (as he rather loosely puts it) 'permanent possibilities of sensation'.²

¹ See Foster (2000: 250–5) and Foster (2008: 107–13). For Leibniz's image, see Leibniz (1712/1989: 199); also Leibniz (1712/2007: 249, 257) and Leibniz (1714/1989: 220).

² See Mill (1865/1889: 187–264). Mill's view comes with a distinctive account of perception, by which a veridical experience isn't one that's caused in the right way, but one that relates to the

To get a sharper definition of phenomenalism, we can start by defining the *mental features* of our world as those that are completely describable in phenomenal terms or a combination of phenomenal and topic-neutral terms. Phenomenal terms are terms for phenomenal properties (of the sort we ascribe to conscious minds and experiences); topic-neutral terms include anything that's fair game for use both in a materialist analysis of the mental and in a phenomenalist (or idealist) analysis of the physical: logical and mathematical terms, terms for various relations of dependence (e.g. causal, counterfactual, and probabilistic), and terms for various modalities (powers, potentials, possibilities, etc.).

The phenomenalism defended here is the conjunction of three claims.

First: conscious experience is irreducible to anything more basic. In this paper, I assume without argument that this claim is correct.³

Second: the physical features of our world supervene on its mental features, in the sense that any possible world indistinguishable from ours in its mental features has all the physical features that our world has. This claim, which I'll call *empirical supervenience*, plays the same role in phenomenalism as psychophysical supervenience plays in materialism.⁴

Third: the mental features on which our world's physical features supervene are pure potentials for conscious experience—pure, in the sense that they aren't metaphysically grounded in anything, and require no explanation in terms of anything except possibly further potentials for experience. This claim, which I'll call *Mill's Thesis*, is one of the things that distinguishes phenomenalism from traditional idealist theories, which locate potentials for experience in the computational architecture or causal powers of some further underlying feature of the world (such as Leibnizian monads, a Berkeleyan God, or Kantian noumena).⁵

$$\text{Phenomenalism} = \left\{ \begin{array}{l} \text{Consciousness Antireductionism} \\ + \\ \text{Empirical Supervenience} \\ + \\ \text{Mill's Thesis} \end{array} \right.$$

totality of all potential experiences in the right way. A discussion of the phenomenalist theory of perception is beyond the scope of this paper, but see Yetter-Chappell (2017) for a closely related idealist account of perception.

³ The arguments against reductionism about consciousness are well known: see Broad (1925), Campbell (1970), Chalmers (1996), Jackson (1982), Kirk (1974), Robinson (1982), and the large literature surrounding these.

⁴ The word 'empirical' comes from the Greek for experience. An clumsier but more revealing label might be 'physicopsychical supervenience'.

⁵ For this reason, phenomenalism is sometimes caricatured as 'Berkeley without God', although a better caricature would be: 'Kant without noumena'.

Why would anyone want to be a phenomenalist? Because phenomenalism has a desirable pair of virtues that no other theory can apparently claim: it's monistic, and it's consistent with a certain sober intuition.

Mind–Body Monism: the mental and physical features of our world are not mutually irreducible.

Sober Intuition: it's possible for a world physically identical to ours to contain no conscious experience.

Many people would like to accept both Mind–Body Monism and Sober Intuition, but few do, since Sober Intuition conflicts with the only kind of monism that most people consider worthy of serious consideration: materialism, the view that the mental features of our world reduce to various physical features of it.

Materialism isn't the only kind of Mind–Body Monism, though, and recent years have seen an uptick of interest in two types of 'mind first' monism: panpsychism, and traditional idealism.⁶

Traditional idealists propose to reduce the physical to the mental by identifying physical phenomena with suitable combinations of conscious experiences: an apple, for instance, consists of the sort of experiences one typically has when one perceives an apple (by sight, touch, smell, taste, etc.).

Panpsychists identify physical phenomena with experiences too, but, unlike traditional idealists, they also identify all experiences with physical phenomena: according to panpsychists, physical states of affairs and phenomenal states of affairs are just the same states of affairs by different names. Like traditional idealists, panpsychists hold that apples are made of experiences, but not the sort of experiences we have when perceiving apples. Rather, apples are made of the experiences (phenomenal character unknown) that panpsychists identify with the apples' constituent atoms.⁷

Panpsychism and traditional idealism are monistic, but they're not consistent with Sober Intuition. If the apples in our world are made of experiences, then it's impossible for a world physically identical to ours not to contain any experience. After all, any world physically identical to ours contains all the apples that our world contains. So, if the apples of our world are combinations

⁶ See, e.g. Freeman (2006) (a collection devoted to contemporary panpsychism) and Goldschmidt & Pearce (2017) (devoted to contemporary idealism). The classic source for traditional idealism is Berkeley (1710/1982), and for panpsychism Eddington (1929).

⁷ One could argue that panpsychism is a kind of materialism (since it equates all mental entities with physical entities) as well as a kind of idealism (since it equates all physical entities with mental entities); see Strawson (2006). Still, panpsychism differs importantly both from traditional materialism (according to which most physical entities aren't mental) and from traditional idealism (according to which many mental entities, such as itches, hallucinations, and dreams, aren't physical).

of conscious experiences, any world physically identical to ours must contain conscious experiences—contrary to Sober Intuition.

Enter phenomenalism.

Like other mind-first metaphysics, phenomenalism proposes to reduce the physical to the mental. However, instead of identifying physical things with experiences, phenomenalists identify them with *potentials* for experience. Potentials for experience aren't experiences, but they still count as mental, provided that we can understand them in purely phenomenal and topic-neutral terms, as phenomenalists hold we can. So, phenomenalism is a kind of Mind–Body Monism.

Most potentials for experience go unrealised in our world, and there is a possible world identical to ours in its potentials for experience, but in which *no* potential for experience gets realised. According to phenomenalism, such a world is physically identical to ours, despite containing no conscious experience. So, phenomenalism is consistent with Sober Intuition.

In short, phenomenalism promises to deliver the Holy Grail of metaphysics: monism without the modal malaise. In this paper, I aim to show that phenomenalism is in a better position to fulfill this promise than people realise.

Section II explains how phenomenalists commit themselves to empirical supervenience by identifying physical phenomena with suitable potentials for experience. Section III defends empirical supervenience from conceivability arguments analogous to those raised against psychophysical supervenience. Section IV elucidates the concept of a pure potential for experience, and explains how such potentials figure in phenomenalist accounts of causation and scientific explanation. Section V defends Mill's Thesis against the objection that ungrounded modalities are ontologically suspect. Section VI concludes the paper.

II. EMPIRICAL SUPERVENIENCE

Phenomenalism is best understood as an identity theory. In this respect, it's analogous to central state materialism (also known as the mind–brain identity theory). This analogy is actually rather instructive; let's look into it further.⁸

Central state materialists hold that the reason why the mental features of our world supervene on its physical features (as they believe) is that the mental features of our world just *are* certain physical features of it, namely brain-states. Central state materialism doesn't identify mental states with isolated brain-states, however. Although materialists sometimes say things like, 'pains are stimulated C-fibers', they're fully aware that if you put some C-fibers in a Petri dish and stimulate them, no pain will result. What pain really is,

⁸ The classic sources for central state materialism are Armstrong (1968) and Lewis (1966).

according to central state materialists, is stimulated C-fibers *suitably integrated with a whole brain*, or at least enough of a brain to support the stimulus-response patterns that materialists consider definitive of pain. Central state materialists see mental states as undetached parts of larger functional wholes.

Analogously, phenomenologists see physical states as undetached parts of larger *mental* wholes. Phenomenologists don't identify physical things with isolated potentials for experience: they're fully aware that a potential for dreaming of a gold brick is insufficient for the existence of a gold brick. What a gold brick really is, according to a phenomenologist, is a potential for experiences as of a gold brick *that cohere with the totality of all potential experiences*.⁹

What does 'cohere' mean, in this context? For an experience to cohere with the totality of all potential experiences is for it to relate to that totality in the way that your present experiences relate to the totality of all the other experiences you've had, as opposed to the way that the experiences you've had in dreams or hallucinations have related to the remainder of your experiences.¹⁰

The physical states that central state materialists identify with mental states are supposed to be categorical features of the world: brain-states, taken as irreducibly non-modal entities. This is the main difference between central state materialism and behaviourism, which identifies mental states with dispositions to respond to stimuli in various ways, and regards the brain-states that underlie such dispositions as explaining, but not being identical with, the mental states.

In this respect, phenomenology is more similar to behaviourism than to central state materialism. Unlike Berkeley, who identifies physical objects with combinations of actual conscious experiences, a phenomenologist identifies them with *potentials* for conscious experiences. Phenomenology is still an identity theory, since it identifies the world's physical features with certain of its mental features. It's just that the mental terms of the phenomenologist identities are potentials for experience, rather than actual experiences.

According to central state materialism, conscious states just are certain physical states; consequently, central state materialism implies that any possible world physically identical to ours contains all the consciousness that our world contains. This is psychophysical supervenience.

According to phenomenology, physical phenomena just are certain potentials for experience; consequently, phenomenology implies that any possible

⁹ Here's a statement of the phenomenologist identity theory that brings out the holistic character of its proposed identifications: every possible world mentally indistinguishable from ours is such that (1) it has all the physical features our world has, and, (2) each of its physical features is identical with some potential for experience. This entails that each physical feature of *our* world is identical with some potential for experience, but the identity is between physical entities and potentials for experience *qua* parts of the totality of all potentials for experience, just as in the mind-brain identity theory, mental states are identified with physical entities (brain states) *qua* parts of totalities of physical states (whole brains).

¹⁰ A full development of phenomenology would replace this working definition of coherence with something more precise; for the purposes of this paper, the working definition should do.

world mentally identical to ours contains all the physical phenomena that our world contains. This is empirical supervenience.

A major objection to materialism is that there are modal counterexamples to psychophysical supervenience. It seems to me that these counterexamples are genuine, and grounds for rejecting materialism. The question naturally arises whether phenomenalism is vulnerable to analogous counterexamples to empirical supervenience. In the next section, I argue that it is not.

III. CONCEIVABILITY ARGUMENTS AGAINST PHENOMENALISM

In this section, we consider three conceivability arguments against empirical supervenience. The first involves a possible world in which all experiences result from interactions between a computer and some envatted brains; I call this the *Matrix Argument*. The second involves a possible world that outwardly resembles our world, but lacks unobservable features (such as elementary particles); this is the *Argument from Unobservables*. The third involves a possible world in which all potential for experience has its basis in disembodied minds; I call this cousin of the Zombie Argument against materialism the *Ghost Argument*.

III.1. *The matrix argument*

The first conceivability argument against empirical supervenience is as follows:

We can conceive of a world in which there hold all the mental facts that hold in our world, but in which those facts hold only because of the operations of a supercomputer connected to some envatted brains. Call this possible world *Matrix*. Any experience or combination of experiences that occurs in our world also occurs in *Matrix*, and any experience or combination of experiences for which there is a potential in our world is an experience or combination for which there's a potential in *Matrix*. However, we can conceive of *Matrix* as being physically very different from our world. For example, we can conceive of it as containing no trees (there's just the computer, the brains, and a barren rocky planet). This gives us a compelling reason to deny that the mental facts about our world (the actual world) logically entail the physical facts about our world.

My response to this argument is to grant the whole thing.

Empirical supervenience says that any possible world that is mentally indistinguishable from ours has all the physical features that our world has. This is different from saying that the mental facts about our world logically entail the physical facts about it. The entailment claim is stronger than the supervenience claim. To show that the entailment claim is false, it's enough to show that there's a possible world that has all our world's mental features, but lacks some of its physical features. To show that the supervenience claim is false, you have to show that there's a possible world that has all *and only* the mental features of our world, but lacks some of our world's physical features.

To see that *Matrix* is not such a world, recall that the mental facts are those that are completely describable using only phenomenal and topic-neutral terms, where topic-neutral terms include any that can legitimately occur both in a materialist analysis of the mental and a phenomenalist (or idealist) analysis of the physical. Although *Matrix* is indistinguishable from our world with respect to the experiences that occur in it, it differs from our world in other mental respects. In *Matrix*, there's a way for there to be experiences as of envatted brains that fails to exist in our world, namely by someone perceiving the brains-in-a-vat setup. This is sufficient for a mental difference between our world and *Matrix*: the idea of 'a way for there to be' is sufficiently portable to count as topic-neutral. (A materialist might equally describe pain as a way for there to be a system satisfying certain functional conditions.)

We can sum the situation up by saying that *Matrix* contains potentials for experience that our world does not: potentials for experiences as of a certain computer-and-envatted-brains setup. Like the idea of a way for there to be something, the idea of a potential is topic-neutral: a materialist may equally speak of a potential for radioactive decay, or a gravitational potential. I'll have more to say about potentials in Section IV; for now, the important point is that in *Matrix*, there are potentials for experience that do not exist in our world (we assume), and that this is a mental difference between our world and *Matrix*.

Can we get around this by modifying the example? Suppose you stipulate a world identical to *Matrix*, except that the supercomputer, brains, and related paraphernalia are for some reason imperceptible—maybe they are shielded from perception by some kind of cloaking device (which also cloaks itself and the rocky planet on which all of this stuff exists), or maybe it's simply a peculiar law of nature that nothing perceives the computer, vats, and so on. Call this scenario *Stealth Matrix*, and the corresponding argument the Stealth Matrix Argument.

Let's concede that there's a sense in which the vat setup is perceptible in *Matrix* but not in *Stealth Matrix*. Still, like *Matrix*, *Stealth Matrix* differs from our world mentally (assuming that we don't live in *Stealth Matrix* ourselves). If what prevents anyone from perceiving the vats in *Stealth Matrix* is a cloaking device, there's still a way for experiences as of vats to occur in *Stealth Matrix* that doesn't exist in our world; namely, through a break-down of the device. If what prevents anyone from perceiving the vats is a natural law, there's still a way for experiences as of vats to occur in *Stealth Matrix* that doesn't exist in our world; namely, through violation of a certain natural law. We might put this by saying that in *Stealth Matrix*, there are second-order potentials for experience that are absent from our world.

It might sound odd to say that there's a way for perceptions of envatted brains to occur in a world in which the laws of nature prevent such perceptions. The important point is that there is a mental *fact*—describe it however you want—that holds in *Stealth Matrix* but not (we assume) in our world: the fact

that certain experiences that might otherwise have occurred fail to occur, due to the existence of a peculiar natural law. In *Stealth Matrix*, there are certain experiences (as of envatted brains) that would occur but for a certain natural law; in our world, this is not the case. This is a mental difference between the two worlds: a difference in a state of affairs fully describable in phenomenal and topic-neutral terms ('experience as of envatted brains', 'natural law', etc.).

The basic challenge for proponents of Matrix-style arguments against phenomenalism is to describe a Matrix scenario in such a way that we can grasp it without thinking of it as differing from the actual world in any mental respect. Rising to the challenge would mean doing what opponents of materialism do when they describe a world physically indistinguishable from ours, but devoid of consciousness. Here, it seems clear what we're being asked to imagine.

By contrast, it's unclear what we're supposed to do, if asked to imagine a world mentally indistinguishable from ours but devoid of trees. When we try, we end up imagining a world that differs from ours in some mental respect, if only by containing potentials for experience that our world doesn't contain.

The first step to mounting a successful conceivability argument is to form a clear conception of a *prima facie* modal counterexample to the target of your argument. The Matrix arguments fail at step one.

III.2. The argument from unobservables

The second conceivability argument against empirical supervenience that I want to consider goes like this:

We can imagine a world observationally indistinguishable from ours, but without any unobservable features. Call it Wysiwyg ('what-you-see-is-what-you-get') World. If there's a potential in our world for certain observations, there's a potential in Wysiwyg World for phenomenally indistinguishable observations, and vice versa. When people in Wysiwyg World visit their counterpart of Niagara Falls, they have the same experiences we have when visiting the actual Niagara Falls; it's just that in Wysiwyg World, the cascading liquid doesn't consist of H₂O molecules or any other microstructure (it's 'Edenic water'). We can stipulate that Wysiwyg World is also indistinguishable from ours in terms of what experiences actually occur in it, as well as in terms of potentials for non-observational experiences (dreams, hallucinations, etc.). Still, since Wysiwyg World lacks the unobservable features of our world (H₂O molecules and so forth), it doesn't have *all* of our world's physical features. The conceivability of Wysiwyg World gives us a compelling reason to deny that the physical features of our world supervene on its mental features.¹¹

Scientists in Wysiwyg World have experiences indistinguishable from those that actual scientists have. For example, they have the same experiences that actual scientists have when using microscopes, cathode ray tubes, Geiger counters, cloud chambers, electrolysis rigs, etc. It follows that the inhabitants of

¹¹ The notion of an Edenic phenomenon comes from Chalmers (2010b).

Wysiwyg World have as much reason as we do to believe that the watery stuff in their world consists of H_2O molecules. It's just that they're wrong (through no fault of their own).

The problem with the argument from unobservables is essentially the same as the one raised earlier for the Matrix arguments. Assuming that the watery stuff that exists in our world does, in fact, consist of H_2O molecules, we have no way to conceive of a world that contains no H_2O molecules but duplicates our world in all mental respects.

Suppose you want to imagine a world W in which there are no H_2O molecules, but in which people nonetheless have exactly the same experiences that actual people (people in our world) have. How do you do it? You could imagine

- (1) that there is some deceiving agent or device in W that gives the inhabitants of W experiences that suggest to them that the watery stuff in their world consists of H_2O molecules, even though it's really Edenic water that has no physical microstructure; or,
- (2) that there are natural laws in W that play the role of the deceiving agent or device described in (1); or,
- (3) that even though the watery stuff in W doesn't consist of H_2O molecules, by a colossal freak-accident people's experiences in W suggest otherwise: scientists always just happen to make certain errors in their calculations, lab equipment always just happens to malfunction in certain ways, etc.

As far as I can tell, these are the only ways to conceive of a world as containing no H_2O molecules despite duplicating our world in terms of what experiences occur in it: by design, by natural law, or by chance. But—and this is the key point—in order to imagine any of these things, we have to imagine a world that differs mentally from ours by containing potentials for experience that our world does not.

In order to imagine the first situation, we have to imagine that certain experiences that don't take place would, were it not for a certain agent or mechanism. In order to imagine the second situation, we have to imagine that certain experiences that don't take place would, were it not for certain natural laws. In order to imagine the third situation, we have to imagine that certain experiences that don't take place would, were it not for a certain statistical fluke. To imagine any such situation is to imagine a world that differs mentally from ours: that is, differs from ours in some phenomenal-cum-topic-neutral respect.

Since the only way to conceive of a world that contains no H_2O is by conceiving of one of the three scenarios described above, and since each of those scenarios involves phenomenal potentials that don't exist in our world, it's impossible to conceive of a world, such as Wysiwyg World was supposed to be, that duplicates our world in all mental respects, but fails to contain H_2O .




Figure 1. Information-equivalent projections of the Earth's surface.

Before moving on to the next conceivability argument, let's briefly consider a different attempt to use unobservables against empirical supervenience.

Suppose we know that one of two empirically equivalent theories is correct, but we don't know which. (By calling the theories 'empirically equivalent', I mean that it's logically impossible for any observation to have different implications for the two theories—e.g. to conflict with one but not the other.) But suppose that despite their empirical equivalence, the theories posit different physical ontologies: one posits zeta particles but no omega waves, the other omega waves but no zeta particles. Then either there's a possible world, *Zeta*, just like ours except that it contains zeta particles instead of omega waves, or there's a possible world, *Omega*, just like ours except that it contains omega waves instead of zeta particles. Since the aforesaid theories are empirically equivalent, both *Zeta* and *Omega* are mentally indistinguishable from our world. The possibility of these worlds would therefore be enough to refute empirical supervenience.

The phenomenalist's best response to this is that empirically equivalent scientific theories are also equivalent in the physical ontologies they posit.

This response conforms to mainstream thinking about how to distinguish between the ontologically significant and the ontologically insignificant differences between different scientific theories. For example, when von Neumann proved that Heisenberg's matrix mechanics was empirically equivalent to Schrödinger's wave mechanics, scientists stopped arguing about which theory was right: they took von Neumann to have shown that matrix and wave mechanics were just different ways of representing the same physical reality.

The idea here is that empirically equivalent scientific theories are like the maps in Fig. 1: they convey the same information in different ways. Naively, one might think that these maps represent different distributions of land and water, but they don't: even though the oceanic projection represents Australia with two non-contiguous shapes, the oceanic projection doesn't say anything about Australia that the continental projection doesn't also say (and vice versa). For some applications, the oceanic projection might be more convenient, for others, the continental projection, but the differences between the two maps are geologically insignificant.

In the same way, the differences between empirically equivalent scientific theories are ontologically insignificant. Like equally accurate projections of the Earth's surface, empirically equivalent theories have the same information value: they differ not in what they say about the physical world, but only in how they say it. Such, at any rate, is the phenomenalist's most natural response to the argument from empirically equivalent theories.¹²

I've argued that we can't conceive of a world that omits some of our world's unobservable physical features without conceiving of a world that differs from ours mentally, at least with respect to the potentials for experience that exist in it. That's not the same as showing how to reduce unobservable physical phenomena to potentials for experience. A fully developed phenomenism would have to carry out such a reduction, at least for all physical unobservables that we have compelling reasons to believe in. Such a reduction is beyond the scope of the present discussion, however, where I've been concerned only to defend phenomenism against the charge that the existence of unobservable physical things entails a failure of empirical supervenience.¹³

III.3. The ghost argument

So far, the conceivability arguments we've considered have all failed, because the hypothetical scenarios on which they relied differed from the actual world mentally, to the extent that they were conceivable at all. The last conceivability argument that I want to consider doesn't suffer from this shortcoming. It is this:

We can conceive of a world that consists of a multitude of disembodied minds. The minds are capable of interaction, and disposed to have various experiences when they do interact. (The minds are powerless to do anything besides cause other minds to have various experiences.) All experiences in this Ghost World arise from such interactions, but not all possible interactions actually take place, and the experiences that occur in Ghost World are only a small subset of those that have the potential to occur there. The experiences that do occur in Ghost World are the same as those that occur in ours, and those that have the unrealised potential to occur in Ghost World (i.e. the ones that have

¹² The principle that empirically equivalent theories have identical ontic import is also known as 'Leibniz equivalence'. In addition to guiding actual scientific practice (as in the case of matrix and wave mechanics), this principle plays a key role in the so-called Hole Argument against spacetime substantivalism: see Earman & Norton (1987) and Norton (1992: 227–30).

¹³ It may be that phenomenism works best in tandem with a limited form of scientific antirealism: it wouldn't be very surprising to learn that the point at which it becomes impossible to phenomenise a scientific posit coincides with the point at which it becomes reasonable to doubt the posit's reality (though not necessarily its conceptual expedience). However, if phenomenists do end up embracing some version of scientific antirealism, it's unlikely to be a version as strong as the one that van Fraassen defends in van Fraassen (1980). According to van Fraassen, we should be agnostic about what hasn't been actually observed by us, whereas the most that a phenomenist would likely have to advocate would be agnosticism about what we can't conceive of being observed by anybody.

the potential to arise from interactions among the disembodied minds) are the same as those for which there is unrealised potential in our world. Ghost World is therefore mentally indistinguishable from ours. But there are no physical objects in Ghost World: it's all just ghostly minds and their experiences. The conceivability of Ghost World gives us a compelling reason to reject empirical supervenience.

My response to this is that there is no physical difference between Ghost World and our world. (So, I agree with the argument up to the part that says that there are no physical objects in Ghost World.)

If there are no physical things in Ghost World, it's not because Ghost World differs from ours in the experiences people have there, or in the experiences that are apt to occur without actually occurring. Ghost World is indistinguishable from ours in all these respects. The only difference between Ghost World and our world is that in Ghost World, but not (we assume) ours, experiences occur because of certain interactions among disembodied minds, and potentials for experience exist because those minds have certain experience-causing powers. So, if there's a reason to doubt that Ghost World contains physical things, it must be that Ghost World, but not ours, fundamentally consists of disembodied minds.

But why should the existence of physical things in a world depend on that world's fundamental constituents having a particular nature, or on their *not* having a mental nature?

Consider an analogy with the history of science. People's beliefs about the ultimate constitution of macroscopic physical objects have changed dramatically over the years, from combinations of the Four Elements, to geometric configurations of Democritean atoms, to dynamical systems of Newtonian bodies, to excitation states of quantum fields. Yet, throughout these changes, people's beliefs about the world's macroscopic physical contents have remained highly stable. The ancient Greeks, the natural philosophers of the Enlightenment, and scientists of the 21st century all agree that the world contains trees, despite having markedly different beliefs about the underlying nature of trees.

Just as different phases in the history of science represent different views about the nature, rather than the existence, of macroscopic physical objects, so different phases in the history of metaphysics represent different views about the nature, but not the existence, of all physical things. A metaphysics that, like panpsychism or Berkeleyan idealism, takes mental entities of some sort as the world's fundamental constituents does not thereby deny the existence of trees or the particles that constitute them. It just offers an unexpected account of their nature.¹⁴

You might raise a semantic objection to the claim that Ghost World contains the same physical objects as our world. You might say that if the experiences

¹⁴ Chalmers defends this position in Chalmers (2010a).

and phenomenal potentials that exist in a given world are grounded in some underlying feature of that world, that feature is a reference magnet for the terms that the inhabitants of that world use. In that case, when someone in Ghost World speaks of a mountain, he refers to something different from anything that we refer to: he refers to disembodied minds, whereas we refer to whatever grounds potentials for coherent experiences of mountains in our world.

To say that *A* grounds *B* is to say, at a minimum, that the existence of *A* both explains and metaphysically necessitates the existence of *B*. The semantic objection fails, because phenomenals deny that anything grounds any potential for experience. Phenomenalists allow that potentials for experience (or at least many of them) have explanations in terms of other potentials for experience—more on this below—but they deny that the explaining potentials metaphysically necessitate the potentials they explain.

In denying that anything grounds phenomenal potential, the phenomenalist isn't being eccentric. Presumably, anyone who rejects reductionism about consciousness will take a skeptical view of the suggestion that phenomenal states or potentials have metaphysical grounds. Only if consciousness reduced to something more basic could there plausibly be a situation in which something both explained and metaphysically necessitated some experience or potential for experience.

Given that potentials for experience have no metaphysical grounds, there are no such grounds for our words to refer to. Rather, our words refer to the potentials themselves, which exist in Ghost World as well as our own. The difference between our world and Ghost World isn't that our world but not Ghost World contains physical things. It's that the existence of physical things has an explanation in Ghost World that it doesn't have in ours: an explanation in terms of a population of disembodied minds.

III.4. Why conceivability arguments against phenomenalism fail

In this section, I've defended phenomenalism against a variety of conceivability arguments. If successful, the defense gives phenomenalism an important advantage over materialism, which is notoriously vulnerable to parallel arguments.

The crucial difference between conceivability arguments against phenomenalism and conceivability arguments against materialism is that the former, but not the latter, rely on demonstrably faulty conceivability claims.

At first, it seems possible to conceive of a modal counterexample to the identification of heat with molecular kinetic energy: just imagine that you have cool sensations when touching an object with high molecular energy. On further consideration, however, we realise that what we've actually conceived of is a hot object that *feels* cool to the touch, i.e. causes phenomenally cool

experiences in those who touch it. That's not the same as conceiving of high molecular energy in the absence of heat.¹⁵

Similarly, it might seem possible, at first, to conceive of a modal counterexample to the identification of physical things with potentials for experience: just imagine some brains in a vat hooked up to a suitably programmed computer, or an Edenic world that has all our world's macrophysical features but none of its microphysical features, or a population of disembodied minds with suitable experience-causing powers. On further consideration, however, we realise that what we've actually conceived of in the first two cases are worlds that duplicate ours at the level of realised experience, but include potentials for experience that don't exist in our world, and, in the third case, a world that differs from ours only in what explains its physical contents, and not in the physical contents themselves.

By contrast, when we conceive of modal counterexamples to the identification of conscious states with brain states, we don't seem to be making the mistake of those who take themselves to conceive of high molecular kinetic energy in the absence of heat. We can, it seems, conceive of people physically just like us who have no experience, or whose visual experiences are colour-inverted relative to ours, and our confidence that we can do so doesn't seem to depend on our overlooking some subtle physical respect in which we've tacitly assumed the imagined people to differ from us.

Conceivability arguments against phenomenalism fail, because they're like conceivability arguments against identifying heat with molecular kinetic energy, and not like conceivability arguments against identifying conscious states with brain states.

IV. MILL'S THESIS

In his original exposition of phenomenalism, Mill introduces the idea of a certain kind of possibility for sensory experience:

The conception I form of the world existing at a given moment, comprises, along with the sensations I am feeling, a countless variety of possibilities of sensation: namely, the whole of those which past observation tells me that I could, under any supposable circumstances, experience at this moment, together with an indefinite and illimitable multitude of others which though I do not know that I could, yet it is possible that I might, experience in circumstances not known to me. These various possibilities are the important thing to me in the world.¹⁶

All physical things are possibilities of sensation, according to Mill, but not all possibilities of sensation are physical things. Define the *phenomenal field* of

¹⁵ The point is Kripke's: see Kripke (1980: 97–155).

¹⁶ Mill (1865/1889: 228).

our world as the hypothetical sum-total of phenomenology that would exist, if all the world's phenomenal potential were realised. In order to count as a physical thing, a possibility of sensation must be a possibility for an experience or combination of experiences that coheres with the other experiences in the phenomenal field, in the sense of 'cohere' explained earlier. If I dream of surfing a mile-high wave, my dream realises a certain potential for experience, but there is no mile-high wave corresponding to that potential, since my dream experience fails to cohere with the rest of the phenomenal field.

Mill calls the experiential possibilities that form the basis of his metaphysics 'permanent possibilities of sensation', 'certified possibilities of sensation', and 'guaranteed possibilities of sensation'. I'm calling them *potentials for experience* or 'phenomenal potentials', for short.

Mill never defines phenomenal potential, although he says enough to make it clear that a potential for experience is more than a mere logical or metaphysical possibility for experience. It's also clear from Mill's remarks that potentials for experience are supposed to be fundamental features of our world, irreducible to anything more basic. Without some further explication, however, the notion of a potential for experience is apt to retain an aura of mystery. Let me say something to dispel this aura.

Take an ordinary example of a potential: a wine glass's potential to shatter. The glass's potential to shatter—its fragility—is grounded in the configuration of the glass's constituent silicon atoms. Fragility is therefore not the kind of potential that phenomenologists are talking about, when they talk about potentials for experience. *Those* potentials aren't supposed to be grounded in anything.

So take a different example: an atom's potential to decay. As far as we know, nothing grounds or explains this potential: its existence is just a basic fact about the atom (or atoms of its kind). This is the kind of potential that Mill's permanent or certified possibilities of sensation are supposed to be. The right model for phenomenalism is not fragility, but radioactivity.

What is it, for there to be a potential for radioactive decay? A sufficient condition seems to be the existence of a non-zero probability for the occurrence of at least one particle-decay event. But this isn't a necessary condition. Alan Hájek discusses the example of an infinitely fine dart thrown at a dartboard with a continuous surface: the dart has the potential to strike the board at a certain point *P*, but the probability that it does strike *P* is zero (one-out-of-infinity). For a more realistic example, if space is continuous, then the probability that a given electron will move to a given point of space at a given moment is likewise zero, even though each point of space is such that the electron has the potential to move there. Likewise, if time is continuous, we can imagine particles with a potential to decay, but whose probability of decaying at any

given moment is the same as the probability of Hájek's dart hitting a given point of the dartboard.¹⁷

Imagine a world of physical objects similar to those that we're familiar with. The objects exist largely in darkness, but some occasionally 'light up', partly or entirely, as if illuminated by an internal or external light-source. We can imagine that when this happens, there is no light source additional to the illuminated object: the object just spontaneously gives off light with the same qualities that would characterise the light that the object would reflect or emit, if an external light source shined on the object from a certain angle, or if a certain part of the object were to start glowing.

Suppose that every object in the imagined world has a potential to light up, though few ever do. Sometimes more than one part of an object lights up at the same time; sometimes a whole object lights up. Some objects are more likely to light up than others, and for some, the probability is zero (like in the dart case). We can also imagine that in some cases, the probability that a certain object will light up in a certain way is tied to the probability that certain other objects will light up in certain ways. Nothing explains why objects have this potential to light up: it's not due to something about their internal structure or anything like that. Illuminability in the imagined world is like radioactivity in ours.

Now replace the illumination events in this example with corresponding experiences—experiences as of viewing variously luminous or illuminated objects—and replace the potentials for illumination events with corresponding potentials for experience. Finally, suppose that all that the world contains are these experiences and potentials for experience.

This is how phenomenalism asks us to think of *our* world. There is a vast, possibly an infinite, number of potentials for experience, some of which get realised, most of which do not. The probability of certain potentials being realised is tied to the probabilities of certain other potentials being realised. Some of the potentials might have only a vanishingly small probability of being realised. The potentials for experience aren't grounded in anything, and, as far as we know, the only thing that ever explains why a potential for experience exists is the existence of some other potential (or potentials) for experience.

The last point calls for elaboration.

Phenomenalists see no need to explain any phenomenal potential in terms of anything besides other phenomenal potentials. But they do require phenomenal potentials, or at least many of them, to have explanations. After all, according to phenomenalists, some phenomenal potentials are physical things, and physical things typically have explanations.

Take an ordinary physical thing, like the Mississippi River delta. The delta is the result of thousands of years of silt- and sand-deposits occurring where

¹⁷ For Hájek's discussion, see Hájek (2003).

the river slows as it enters the Gulf of Mexico. Like anyone else, a phenomenalist recognises that the delta is a natural consequence of these hydrological processes. It's just that a phenomenalist sees both the delta and the processes that created it as potentials for experience.

The motions of water and sediment reduce to certain phenomenal potentials, the delta reduces to certain other phenomenal potentials, and the existence of the latter potentials is a non-metaphysical (causal, natural, or nomological) consequence of the existence of the former. The phenomenal potentials that constitute the hydrological processes naturally necessitate the phenomenal potentials that constitute the delta.¹⁸

In short, phenomenologists hold that many (perhaps all) potentials for experience have non-reductive explanations *in terms of other potentials for experience*. If there's an established scientific explanation for why a certain potential for experience exists, phenomenologists can happily accept it. They merely add that the terms of the scientific explanation are themselves potentials for experience.¹⁹

This also goes for neuroscientific explanations.

Phenomenalism treats brains the same way it treats other physical things: as potentials for experience. Your brain, for example, is a potential for experiences like those that we'd have if we were observing your brain (while performing brain surgery on you, or giving you an MRI scan, or whatever). Phenomenalism accounts for the physical effects of brains the same way it accounts for the physical effects of other things: the relationship between your brain-activity and your motor behaviour is the same as that between the motions of water-borne sediment and the delta.

However, brains don't have only physical effects: brain-activity also causes, or at least correlates with, the experiences that make up conscious mental lives. Furthermore, it often happens that there are correlations between experiences associated with different brains. How does phenomenalism account for all this?

Suppose we're playing catch with a baseball. As we play, there's a correlation between your visual experiences and mine: as my visual impressions of the ball

¹⁸ As Mill puts it, 'Whether we are asleep or awake the fire goes out, and puts an end to one particular possibility of warmth and light. Whether we are present or absent the corn ripens, and brings a new possibility of food. Hence we speedily learn to think of Nature as made up solely of these groups of possibilities, and the active force of Nature as manifested in the modification of some of these by others'. Mill (1865/1889: 230) See also Ayer (1940: 229–31) and Ayer (1946–1947: 146–50).

¹⁹ Phenomenalism is neutral on whether every potential for experience has an explanation (in the form of further potentials for experience). In this, phenomenalism is no different from materialism, which is neutral on whether every physical state has an explanation (in terms of further physical states). Just as materialism is compatible with the existence of inexplicable physical states, phenomenalism is compatible with the existence of inexplicable potentials for experience.

shrink (occupy less of my visual field), your visual impressions of the ball correspondingly grow, and vice versa. What explains this correlation?

Roughly, the explanation is as follows. A baseball is moving back and forth between two hominids, each equipped with eyes and brains that function in certain ways. The amount of each hominid's retinal surface that's stimulated by light reflected from the ball is proportional to the distance between the ball and his eyes; since this distance is inversely correlated for each hominid (when the ball is near you, it's far from me, and vice versa), the amounts of retinal stimulation are also inversely correlated. Due to the way our brains are organized and connected to our eyes, the inversely sized retinal stimulations cause the visual centres of our brains to go into neural states that are also inversely correlated in terms of some relevant magnitude (e.g. the number of neurons activated in our visual cortices' retinotopic maps). Since there is a lawlike correlation between the occurrence of such neural states and visual experiences of baseballs—that is, since those brain states are the neural correlates of such experiences—your baseball-impressions grow as mine shrink, and vice versa.

A phenomenalist can accept this explanation. *Of course* we have bodies; *of course* there's a baseball moving back and forth between them; *of course* light from the ball is interacting with our eyes in various ways, resulting in various patterns of brain activity which take place when, and only when, visual impressions of baseballs occur. Phenomenalism is compatible with all of this. It's just that, according to the phenomenalist, facts about bodies, balls, light, etc. reduce to mental facts—facts about phenomenal potentials.

The correlation between our visual experiences is a mental fact that has an explanation that invokes various physical conditions (involving our bodies, light, and the ball), together with an empirical principle to the effect that certain kinds of brain-activity correlate with certain forms of experience. This is not a reductive explanation, of course; for a phenomenalist, as for anyone who opposes reductionism about consciousness, a reductive explanation of correlations among different subjects' conscious experiences is as impossible as a reductive explanation of consciousness itself. But it is an explanation, and one that's consistent with the phenomenalist position, provided that we understand the physical conditions as potentials for experience.

What is the status of the empirical law that certain kinds of brain-activity correlate with certain forms of experience? According to phenomenologists, this correlation is not a case of identity: like others who reject materialism, phenomenologists hold that it's possible for the neural correlates of our experiences (i.e. by phenomenalist lights, certain phenomenal potentials) to exist in a world in which there is no experience. The correlation between brain-activity and experience is at most nomologically necessary, not logically or metaphysically necessary.

An implication of this is that phenomenologists, like dualists, are under some pressure to deny that our conscious experiences cause any of our bodily

behaviour. Given that we can already account for our behaviour as arising from purely neural causes, it's not easy to see how our conscious experiences—pains, visual images of stop signs, auditory impressions of alarm clocks, etc.—could plausibly influence our behaviour.²⁰

Does this mean that phenomenalism works only in conjunction with epiphenomenalism? Maybe. But three qualifications are in order.²¹

First, any move that a dualist can make to avoid epiphenomenalism is also available to a phenomenalist. For example, if, as some dualists argue, an overdeterministic version of interactionism is defensible, an epiphenomenalist can subscribe to that kind of interactionism.²²

Second, any measure that a dualist can take to make epiphenomenalism tolerable is also available to a phenomenalist. For example, phenomenologists, like dualists, can point out that all of the observed correlations between mental events and physical events are compatible with a causal story in which conscious experiences don't cause bodily behaviour, but in which the experiences and behaviour are effects of a common physical cause. Phenomenologists can also co-opt dualist explanations of how we know about our own conscious states, despite those states being physically inefficacious.²³

Third, whereas epiphenomenalist dualism makes conscious experience completely irrelevant to the physical world, this isn't true of epiphenomenalist phenomenalism: here, conscious experience *is* relevant, as the realiser of that which constitutes physical reality. Let me elaborate.

In the dualist view, there's no deep metaphysical connection between consciousness and the physical world. There are robust correlations between the physical processes that occur in our brains, and the conscious processes that occur in our minds, but that's the whole extent of the connection between mind and matter, according to dualism. To explain the correlations, or their robustness, dualists may posit natural laws requiring certain conscious states to occur whenever certain brain states occur; but, if dualism is true, the existence of such laws is no less surprising or inexplicable than the existence of consciousness itself. From a dualist standpoint, nothing about the physical world gives us any reason to expect there to be such a thing as conscious experience at all: consciousness enters the picture as something completely new and beyond anything one could reasonably expect on the basis of physical information alone.

With phenomenalism, the situation is different. There is a deep metaphysical connection between consciousness and the physical world: the physical

²⁰ See Kirk (2005) for a detailed discussion of this point.

²¹ I understand epiphenomenalism as the view that conscious experiences have physical causes, but no physical effects.

²² For a defense of overdeterministic interactionism, see Mills (1996).

²³ For the common-effects account, see Jackson (1982) and for an account of our knowledge of physically inefficacious experience, Chalmers (1996: 172–209).

world consists of potentials for conscious experience. Admittedly, the existence of these potentials doesn't guarantee that there is any actual experience; in this sense, consciousness remains a mystery. However, if physical things are potentials for experience, then the actual occurrence of conscious experiences is less surprising than it otherwise would be, since all it requires is for some of those potentials to be realised. From a phenomenalist standpoint, there *is* something about the physical world that gives us a reason to expect there to be conscious experience. Consciousness enters the picture as something new, yes, and even to some extent unexpected, but not *as* unexpected as in the dualist scheme of things.²⁴

V. PHENOMENALISM VS THE ARMSTRONG DOCTRINE

The idea that the physical world is phenomenal potentials all the way down conflicts with an influential doctrine of late 20th century metaphysics. David Armstrong was probably the doctrine's leading proponent, insisting throughout his long career that you can't just have free-floating powers: something must *have* the powers, and this something can't just be more powers. More generally, Armstrong held that there is no such thing as primitive modality: any powers, potentials, dispositions, or possibilities that exist in our world must exist by virtue of our world's having some non-modal features. Let's call this the *Armstrong Doctrine*.²⁵

The Armstrong Doctrine is incompatible with Mill's Thesis. Consequently, if the doctrine were true, we'd have to abandon phenomenism. Fortunately for phenomenists, the debate over whether powers, potentials, and similar modalities require categorical grounds has gone rather strongly against the Armstrong Doctrine in recent decades.²⁶

Consider radioactivity again. Specifically, consider radon atoms. These have a potential to undergo radioactive decay: there's about a 50% chance that a radon atom decays within a four day period.²⁷

Presumably, about 50% of the radon atoms that come into existence in our world decay within four days. But we can imagine a world categorically

²⁴ It's true that not all experiences realise potentials that partly constitute physical things, according to phenomenism: the conscious states that occur in us during dreams, hallucinations and the like do not. Such states realise potentials for experiences that fail to cohere with the total field of potential experiences.

²⁵ See Armstrong (1961: 56–58) and Armstrong (1993: 187); also Lewis (1966: 20), Lewis (1992: 218–19), and Lewis (1998).

²⁶ The world's categorical features are those it has other than by virtue of the existence of some power, potential, disposition, or related modality. According to phenomenism, our world's only categorical features are the conscious experiences it contains.

²⁷ More accurately, there's about a 50% chance that the quantum tunnelling involved in the decay of a radon-222 atom occurs within four days of the atom's coming into existence.

indistinguishable from ours, in which there's a 90% chance that a radon atom decays within four days. We need only imagine that due to a colossal statistical fluke, about 50% of the radon atoms in this other possible world decay within four days, despite there being a 90% chance of any given radon atom decaying within four days. In this other world, there exist potentials for radioactive decay that do not exist in our world, despite the other world's being categorically indistinguishable from ours (duplicating our world with respect to its non-modal features). It follows that the potentials for radon decay that exist in our world don't exist due to our world's having some categorical feature or features.²⁸

Furthermore, it is, as far as we know, a physically fundamental fact about radon atoms that they have about a 50% chance of decaying within four days. We have no reason to think that this fact has an explanation in terms of some non-modal feature of radon atoms. If you like, you can say that the potential for decay inheres in the atoms, but the potential that thus inheres is, as far as we know, irreducibly modal. This is so, even if we assume that atoms themselves are categorical features of our world, and not ungrounded potentials of some sort. On that assumption, the fragility of a wine glass—its potential to shatter—arguably does reduce to certain categorical features of the glass (atoms arranged in a certain shape, and held together by certain bonds). But a radon atom's potential to decay isn't like this, at least not according to our best science.²⁹

So, not only is it metaphysically possible for a potential to exist without any categorical basis: we also have reason to believe, or at least not to disbelieve, that many of the potentials that actually exist have no categorical basis. It follows that when phenomenologists say that potentials for experience lack any basis (i.e., that they aren't grounded in anything, categorical or otherwise), no one can accuse them of positing a kind of thing that we would otherwise have no reason to admit into our ontology.

We phenomenologists deny that potentials for experience are grounded in or reducible to anything, and we affirm that many (possibly all) phenomenal potentials are explained by other phenomenal potentials. But we need neither affirm nor deny that phenomenal potentials have an explanation in terms of entities that aren't phenomenal potentials (such as monads, God, or noumena). It's unclear what could justify positing such entities, but as phenomenologists we can afford to be agnostic about this. We only insist that if the phenomenal potentials of our world do have some explanation in terms of entities that are not themselves phenomenal potentials, physical things are still to be identified

²⁸ This is the central argument of McKittrick (2003), as I understand it. The focus of McKittrick's discussion is dispositions, but the points she makes also apply to potentials.

²⁹ The argument here is basically the one that Stephen Mumford gives for the reality of ungrounded dispositions: see Mumford (2006).

not with those explanatory entities, but with the phenomenal potentials whose existence they explain.

The Great Sphinx of Giza is an ungrounded potential for experience, or an assemblage of such potentials. Modally, the Sphinx goes wherever the potentials go. If they go to a world where their existence has some deeper explanation (in terms of a population of noumena or disembodied minds, for example) the Sphinx goes there too. If they go to a world in which their existence has no deeper explanation, the Sphinx follows. But phenomenologists don't think that the potentials that constitute the Sphinx necessarily have to go anywhere in order to exist in a world in which their existence has no deeper explanation—no explanation, that is, in terms of anything but further potentials for experience. As far as we know, they exist in such a world already.

Since phenomenology doesn't ground potentials for experience in actual experiences, one might wonder whether it's really a kind of monism, as advertised. It's true that phenomenology posits two, mutually irreducible features of the world: experiences, and potentials for experience. But this doesn't make phenomenology into a kind of dualism. Materialists who recognize an irreducible distinction between physical events and certain potentials for physical events, such as potentials for radioactive decay, are not on that account classified as dualists. No more should we classify phenomenologists as dualists for recognising an irreducible distinction between experiences and potentials for experience.

VI. CONCLUSION

A lot has happened in metaphysics since phenomenology last had serious defenders. For one, we now have a much better understanding of how to evaluate identity and necessity claims; for another, we have a much fuller picture of powers, dispositions, and related modalities. Since these developments are directly relevant to phenomenology, it seemed prudent to consider whether they have altered its prospects in any way.³⁰

I've argued that they've brightened its prospects considerably: phenomenology is immune to the sort of modal counterexamples that beleaguer material-

³⁰ Fumerton dates phenomenology's demise to Chisholm (1948), which argues, in effect, that Mill's Thesis is incompatible with empirical supervenience: according to Chisholm, no proposition asserting the existence of a potential for experience is equivalent to any proposition asserting the existence of a physical state of affairs. (For Fumerton's discussion, see Fumerton [1985: 141-5].) Chisholm's target is the phenomenology of Lewis (1946: 203-53) and Ayer (1946-1947), according to which the existence of a phenomenal potential reduces to the truth of a conditional of the form: 'If such-and-such phenomenal conditions were satisfied, such-and-such other phenomenal conditions would be satisfied'. Since the phenomenologist theory I've defended here doesn't construe phenomenal potentials this way, it is immune to Chisholm's criticism.

ism, and the Armstrong Doctrine, once considered the most serious challenge to phenomenalism, is now regarded as doubtful at best.

One would have to say much more to restore phenomenalism to the prominence it once enjoyed, if such a restoration is even possible, or desirable. My purpose has been the humbler one of defending phenomenalism against some of the more important and influential objections to it. The hope is that by overcoming those objections, we can return phenomenalism, if not to its former glory, at least to the metaphysics syllabus.³¹

REFERENCES

- Armstrong, D. (1961) *Perception and the Physical World*. London: Routledge & Kegan Paul.
- (1968) *A Materialist Theory of the Mind*. London: Routledge & Kegan Paul.
- (1993) 'Reply to Martin', John Bacon, Keith Campbell and Reinhardt Lloyd (eds), *Ontology, Causality, and Mind: Essays in Honour of D.M. Armstrong*, 186–194. Cambridge: Cambridge University Press.
- Ayer, A. J. (1940) *The Foundations of Empirical Knowledge*. London: Macmillan.
- (1946–1947) 'Phenomenalism', *Proceedings of the Aristotelian Society*, 47, 163–96.
- Berkeley, G. (1710/1982) *A Treatise Concerning the Principles of Human Knowledge*. K. Winkler, ed. Indianapolis: Hackett Publishing.
- Broad, C. D. (1925) *Mind and Its Place in Nature*. New York: Harcourt, Brace & Co.
- Campbell, K. (1970) *Body and Mind*. London and Toronto: Macmillan.
- Chalmers, D. (1996) *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- (2010a) The Matrix as metaphysics. *The Character of Consciousness*, 455–494. Oxford: Oxford University Press.
- (2010b) Perception and the fall from Eden. *The Character of Consciousness*, 381–454. Oxford: Oxford University Press.
- Chisholm, R. (1948) 'The problem of empiricism'. *The Journal of Philosophy*, 45(19), 512–17.
- Earman, J. and Norton, J. (1987) What price spacetime substantivalism? The hole story. *British Journal for the Philosophy of Science*, 38(4), 515–525.
- Eddington, A. S. (1929) *The Nature of the Physical World*. New York: Macmillan.
- Foster, J. (2000) *The Nature of Perception*. Oxford: Oxford University Press.
- (2008) *A World for Us: The Case for Phenomenalistic Idealism*. Oxford: Oxford University Press.
- Freeman, A. (ed.) (2006) *Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism?* Exeter: Imprint Academic.
- Fumerton, R. A. (1985) *Metaphysical and Epistemological Problems of Perception*. Lincoln, Nebraska: University of Nebraska Press.
- Goldschmidt, T. and Pearce, K. L. (eds). (2017) 'Idealism: New Essays in Metaphysics'. Oxford: Oxford University Press.
- Hájek, A. (2003) 'What conditional probability could not be'. *Synthese*, 137(3), 273–323.
- Jackson, F. (1982) 'Epiphenomenal qualia'. *The Philosophical Quarterly*, 32(127), 23–36.
- Kirk, R. (1974) 'Sentience and behavior'. *Mind*, 83(329), 43–60.
- (2005) *Zombies and Consciousness*. Oxford: Clarendon Press.
- Kripke, S. (1980) *Naming and Necessity* (revised edition). Oxford: Basil Blackwell.

³¹ Many thanks to Zach Barnett, Bob Beddor, Ben Blumson, Brian Cutter, Thomas Hofweber, David Papineau, Abelard Podgorski, Qu Hsueh Ming, Neil Sinhababu, Rob Smithson, the participants of the 2017 NYU-Shanghai Idealism Workshop, and two anonymous referees for their sharp and constructive criticism.

- Leibniz, G.W. (1712/1989) Notes for Leibniz to Des Bosses, 5 February 1712. Roger Ariew and Daniel Garber (eds), *Philosophical Essays*, 199–200. Indianapolis and Cambridge: Hackett Publishing Company.
- (1712/2007). *Leibniz to Des Bosses, 16 June 1712*. Brandon C. Look and Donald Rutherford (eds), *The Leibniz-Des Bosses Correspondence*, 255–257. New Haven: Yale University Press.
- (1714/1989). ‘Monadology’. Roger Ariew and Daniel Garber (eds), *Philosophical Essays*, 213–225. Indianapolis and Cambridge: Hackett Publishing Company.
- Lewis, C. I. (1946) *An Analysis of Knowledge and Valuation*. La Salle: Open Court.
- Lewis, D. (1966) ‘An argument for the identity theory’. *Journal of Philosophy*, 63(1), 17–25.
- (1992) ‘Critical notice of *A Combinatorial Theory of Possibility*’. *Australasian Journal of Philosophy*, 70(2), 211–24.
- (1998) ‘The truthmakers’. *Times Literary Supplement*, 4948, 30.
- McKittrick, J. (2003) ‘The bare metaphysical possibility of bare dispositions’. *Philosophy and Phenomenological Research*, 66(2), 349–69.
- Mill, J. S. (1865/1889) ‘An Examination of Sir William Hamilton’s Philosophy, and of the Principal Philosophical Questions Discussed in his Writings’. Sixth edition. London: Longman, Green, and Co.
- Mills, E. (1996) ‘Interactionism and overdetermination’. *American Philosophical Quarterly*, 33(1), 105–17.
- Mumford, S. (2006) ‘The ungrounded argument’. *Synthese*, 149(3), 471–89.
- Norton, J. D. (1992) ‘Philosophy of space and time’. Merrilee H. Salmon (ed), *Introduction to the Philosophy of Science*, 179–231. Englewood Cliffs: Prentice Hall.
- Robinson, H. (1982) *Matter and Sense: A Critique of Contemporary Materialism*. London: Cambridge University Press.
- Strawson, G. (2006) ‘Realistic monism: why physicalism entails panpsychism’. Anthony Freeman (ed), *Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism?* 3–29. Exeter: Imprint Academic.
- van Fraassen, B. (1980) *The Scientific Image*. Oxford: Oxford University Press.
- Yetter-Chappell, H. (2017) ‘Idealism without God’. Tyron Goldschmidt and Kenneth L. Pearce (eds), *Idealism: New Essays in Metaphysics*, 66–81. Oxford: Oxford University Press.

National University of Singapore, Singapore