Penultimate Draft

Final version appears in:

Imagination and Experience: Philosophical Explorations, Vendrell Ferran & Werner, eds., Routledge, 2024.

The Unimaginability of Experience

Peter Langland-Hassan

1. Light metaphors

Light metaphors abound in the characterization of consciousness. Philosophers describe consciousness as a "mysterious flame" (McGinn 2000), noting that, in its absence, "all is dark inside" (Chalmers 1996, 96). A central task for philosophical theories of consciousness, it is said, is to explain how conscious states "light up" (Rosenthal 2004, 19). Neuroscientists describe acquiring consciousness as "stepping into the light" and identify the question of how there comes to be a "movie in the brain" as one of the chief stumbling-blocks to the scientific explanation of consciousness (Damasio 1999, 4, 8). Information-processing in the absence of consciousness is said to occur wholly "in the dark" (Koch 2004, 231). Such metaphors are especially common in characterizations of the kind of consciousness that seems to elude capture by traditional modes of explanation—the "phenomenal consciousness" (Block 1995, 230-231) that generates the "hard problem" of consciousness (Chalmers, 1996).

We might ask *why* light metaphors seem especially appropriate to this task. Consider a (fictional) robot, Blink, that is navigating its environment, discriminating among shapes and colors, heading cross-country, to California, with the aid of GPS navigation. It seems intuitive to think of his cognition—his information processing—as occurring wholly "in the dark." Yet, if we imagine the cognitive circuitry inside Blink's cranial chassis as inexplicably suffused in light, as pulsating bright blues, greens, and reds with intermittent flashes of bright white, it becomes just a bit more plausible to think that Blink has somehow acquired consciousness. Of course, we

know better than to think that Blink is conscious just due to the lights glowing inside his cranium. Nevertheless, a filmmaker's intention that Blink be viewed as conscious would be clear if they represented the robot in that way.

Why should that be the case? Why should light metaphors track intuitions about consciousness? Why does their metaphorical use seem less arbitrary than would, say, heat- or speed-related metaphors? A deflationary answer is that there is nothing especially fitting about light metaphors and that their intuitive appropriateness is simply ingrained through past experience. On this view, heat or speed metaphors *could have* been just as evocative of consciousness, had there been a similarly longstanding practice involving their use. This is, I think, the wrong answer. I will sketch what I see as the right answer now and will spend the remainder of this chapter defending it.

People think of the mind as being well-lit, metaphorically speaking, because they take themselves to imagine experiences. By 'experiences', I do not mean sequences of worldly events, such as glasses falling off of tables, but, rather, *conscious mental events* of the sort that present the world as being a certain way, and which can occur when the events they represent do not.[1] When we (try to) imagine experiences of this sort, we (often) use visual imagery. And, like visual perceptual experiences, visual images typically represent well-lit environments; they do not, after all, represent unilluminated, non-visible colors and objects. Visual imagery represents colors and objects that are more or less suffused in light. Given that we form visual imagery while thinking about experiences—our own, and those of others—we are left with the nagging sense that they are somehow well-lit or illuminated. Like "tropical sea-water" they appear "phosphorescent" and "self-luminous" (Ryle 1949/2009, 159). A puzzle then emerges for materialist theories of mind: how could it be that brain processes are quasi-illuminated in this way? How could they have these *phenomenal properties*?

The question rests on a mistake. When we think we are imagining experiences—and thereby representing them through the use of visual imagery—we are actually *mis*imagining them in a

[1] It is common in analytic philosophy of mind to use the term 'experience' to refer to conscious mental entities of a kind. This may sound like an overly narrow notion of "experience" for those working in a phenomenological tradition. It loses nothing in what follows to substitute "conscious mental states" for "experiences."

certain way. Experiences *themselves* are not directly imaginable. [2] It is the conviction that they are so imaginable that leads to the view that they have luminous properties that cannot be squared with the neuroscientific account of such states. Is the same conviction that underlies the seeming appropriateness of light metaphors.

In speaking of what we can or cannot directly imagine, I am interested in the limits of sensory (or "perceptual") imagination. Sensory imagining necessarily involves mental imagery (Van Leeuwen 2013; Langland-Hassan 2020; Kind 2001). Whether all, or only some, uses of mental imagery are instances of sensory imagining can be left open, for our purposes. The most salient form of mental imagery for humans is visual imagery, as is triggered when one visualizes a scene; but there are plausibly instances of sensory imagining across the sense modalities (e.g. "hearing a song in your head"). My thesis is that we cannot *sensorily* imagine experiences in the direct, unmediated way in which we can sensorily imagine ordinary objects like tomatoes, bicycles, and chairs.

Precisely what it means to be *directly* imaginable will be sorted out below. The answer will entail that experiences are not imaginable in anything like the way many take them to be, and, hence, that we do not have the kind of (phenomenal) knowledge of experiences that many think we have. This chapter can thus be considered an entry in the growing literature on illusionism—the view that intuitions concerning the inexplicability of consciousness by ordinary scientific means are founded on a kind of illusion to which we are naturally susceptible when thinking about consciousness (Frankish 2016). I will begin, next section, by presenting defeasible evidence that sensory imaginings directly represent the kinds of things that perceptual experiences directly represent: namely, the surface properties of worldly objects. Then, in Section Three, I counter Christopher Peacocke's (1985) claim that all imagining requires (directly) imagining an experience and propose a better-motivated framework where experiences are only indirectly imaginable. Section Four draws on that framework to give a diagnosis of why

[2] One might agree with this point with respect to visual imaginings (that make use of visual imagery), while holding that there are other forms of imagery—such as imagery of bodily pains and pleasures—that can be used to veridically (and directly) imagine their corresponding mental states. However, on a representationalist perceptual account of pain (Dretske 2003; Aydede 2019), what goes for visual experiences and visual imagery goes for bodily sensations and their corresponding imagery as well—as further discussed in Section 2.

light metaphors gain traction in discourse about the mind and how they lead to errors. Section Five concludes by rebutting an argument from M.G.F. Martin (2002) that imagining ordinary objects always requires one to (directly) imagine an experience.

2. Perceiving and imagining

We can imagine what we can see. I just saw Bob; now I am imagining Bob. We cannot see experiences (not without someone's skull being opened). Does sensory imagination *outstrip* visual perception in taking experiences as its objects as well? Here are some *prima facie* reasons for thinking it doesn't:

1) Neuroimaging studies have long revealed significant (though not complete[3]) overlap in the regions activated during vision and visual imagery (Kosslyn et al. 1999; Ganis, Thompson, and Kosslyn 2004). Recent research shows that such similarities are *content-specific*. Stokes et al. (2009) were able to generalize from patters of activity in high-level visual cortex during acts of perceiving X's and O's to predict (via fMRI) whether a person was imagining an X or an O. Others have shown that areas activated during the visual perceptual representation of faces, houses, and places, respectively, are also activated during visual imagery tasks with the same contents (O'Craven and Kanwisher 2000; Mechelli et al. 2004).

2) Performance on visual imagery tasks degrades when subjects are given a concurrent visual perception task (Craver-Lemley and Reeves 1992; Brooks 1967); the same pattern of interference does not occur when the concurrent task is not visual in nature.

3) Stephen Kosslyn has reported many studies where the performance time required for visual perception tasks is mirrored by the time required for corresponding visual imagery tasks (see Kosslyn (1994, p. 4-10) for a summary). In an example of one such study, partly aimed at controlling for criticisms levied by Pylyshyn (2002, 2003), Borst and Kosslyn (2008) found that

[3] The lack of complete overlap is often hypothesized to result from perceptual representations requiring "bottom-up" up activity (from early visual areas), while corresponding imagery representations are triggered by "top-down" signals from pre-frontal cortex (Mechelli et al. 2004; Ganis, Thompson, and Kosslyn 2004). The basic idea is that the brain has two ways of triggering the same processes. Such an interpretation is compatible with the double-dissociations between visual perceptual and visual imagery abilities noted by Bartolomeo (2008).

subjects took proportionately longer to scan proportionately longer distances between two points on a grid, both when the grid was visually perceived and imagined.

4) Visual perceptual deficits are often mirrored in one's capacity to use visual imagery. For instance, patients who, due to brain lesions, neglect a certain region of the visual field will neglect that same region in visual imagination (Bisiach, Luzzatti, and Perani (1979)).[4] Visual agnosias—e.g., the inability to recognize objects as being of a certain kind by sight—have also been seen to be mirrored in patients' imagery (Levine, Warach, and Farah 1985).

Taken together, these neurological and functional similarities between visual imagery and perception support the presumption that the informational limits of each—in terms of the *kinds of things* they can represent and thereby provide information about—will be roughly the same. We need a good reason to abandon the symmetry and hold that sensory imagination can take as its direct objects *both* experiences and everyday objects, while vision can only be directed at the latter.

Yet, as earlier noted (fn. 2), one might accept this point for visual imagery, while maintaining that we have imagery of bodily sensations (such as pains and pleasures) that *can* be used to directly imagine those experiences. If that were the case, then at least some "experiences" (e.g., pain experiences) could be sensorily imagined. However, the same points apply to pains and pain imagery as apply to visual perception and visual imagery, if we adopt a perceptual account of pain experience (Dretske 2003; Langland-Hassan 2017; Byrne 2001).[5] Suppose that you hit your thumb with a hammer. The resulting pain experience represents (painful) tissue damage as occurring in a specific place: your thumb. There are, in fact, many kinds of bodily damage and distress represented in correspondingly different ways by different kinds of pain experience (sharp, dull, throbbing, hot, prickly, and so on). The variety of pains, considered as mental states, represent and thereby provide information about these unwelcome bodily states. This, indeed, is the cognitive role and evolutionary rationale of pains. Just as visual imagery provides information about the local external environment, bodily sensations (like pains) provide

[4] Such findings extend to non-visual motor imagery as well. Sirigu *et al.* (1995) found that deficits in a patient's ability to move her fingers were attended by similar deficits in her ability to imagine moving her fingers.
[5] Whether a representationalist perceptual theory of pains and other bodily sensations is correct is a matter of debate. See Aydede (2019) for an overview.

information about one's body. All the same, neither represent or provide information about *themselves.*

It should, at least, be our presumption that the representational limits of visual and pain imagery are (at least roughly) the corresponding limits of visual and pain perceptual states. We need good reason to abandon the symmetry. What sort of reasons have people offered?

3. Imagining an experience as the common factor?

Christopher Peacocke finds it "uncontroversial" that, when it comes to sensory imagination, "to imagine being X…is always at least to imagine from the inside an experience as of being X" (1985, p. 21-22). Not only does this hold for instances of imagining *being* X, says Peacocke, it also holds "for such forms as 'imagining a valley'," where "we can say that to imagine an F is always at least to imagine from the inside an experience as of an F" (p. 23).

For Peacocke, imagining a valley requires "at least" that one imagine "from the inside, an experience as of" a valley. He does not deny that we can imagine objects; he just holds that doing so requires imagining an experience. How do we manage to *simultaneously* imagine objects and experiences? Peacocke distinguishes between the mental images used in an act of imagining and the conditions that are "S-imagined" to hold in the situation (where, he explains, 'S' is for 'suppose'). S-imagining "shares with supposition the property that what is S-imagined is not determined by the subject's images, his imagined experiences" (p. 25). Notice that Peacocke glosses "the subject's images" as "his imagined experiences." This indicates that, for Peacocke, the use of a mental image by itself suffices for an experience to be imagined. Various S-imagined conditions may then make it true that an object is imagined *in addition*.

However, Peacocke does not provide an argument that simply having an image guarantees that an experience is imagined. He takes it as obvious, a "conceptual truth." So there is no straightforward rebutting to be done. Yet a kind of argument for this would-be conceptual truth can be extracted from the work he sets out for images and S-imagined conditions. He observes that the same mental image can serve in imaginative projects where different scenarios are imagined. To take Peacocke's example, the same image can (in some sense) be featured in an instance of imagining a suitcase, imagining a suitcase wholly obscuring a cat, and even in

imagining yourself hallucinating a suitcase. How can this be so? In particular, how can it be that one type of image is suitable *both* for imagining an experience—a mental process—and an ordinary suitcase? These are rather different things! Peacocke's answer is that the same image can serve in such cases because "the imagined experience fulfilling each project can be the same" (1985, p. 24). The common factor in these cases is the same kind of experience is imagined; differences only emerge at the level of whatever else is "S-imagined" to be true in the situation. The answer to the puzzle of how the same type of image can be appropriate to imagining an experience and an object is that, in each case, the image itself is of the same thing: a certain type of experience, common to each imagined scenario.

On this view, there is an important asymmetry: we can imagine a cat behind a suitcase *by means of* sensorily imagining a certain kind of experience; but we cannot imagine an experience by means of sensorily imagining a cat. Imagining a cat (or any worldly object) requires that we exercise an ability to imagine experiences, and not vice versa. Noting this asymmetry is one way of capturing the intuitive distinction between *directly* sensorily imagining something and sensorily imagining it in some other, indirect sense. For Peacocke, it can be said that we directly imagine experiences—and not cats—because every sensory imagining must always be an imagining of an experience, whereas not all sensory imaginings must be imaginings of non-mental things in the world (as when, for instance, we imagine a hallucination). We indirectly imagine suitcases by directly imagining experiences as of suitcases and then adding, as a kind of attached supposition (an "S-imagining"), that there is a cat hiding behind the suitcase.

Whatever its intuitive merit, this way of viewing things clashes badly with the earlier-noted parallels between sensory imagining and perceiving. Perceiving an *x* is not always (or perhaps ever) to perceive an experience as of an *x*. We *have* experiences and by means of them perceive things in the world. But we do not need to perceive our experiences themselves in order to perceive objects. The purpose of perception is not to make us aware of our own minds, after all, but to allow awareness of our immediate environment. Considering, also, the profound neurological and functional similarities between sensory imagining and perception reviewed in Section Two, it would be surprising if sensory imaginings were—by nature—things that represent (or make us aware of) *experiences*, while perceptual experiences themselves were not.

Such a discontinuity remains possible, of course. But we have reason to explore whether there are other ways of understanding the common factor in cases of the kind Peacocke discusses that do not require positing a fundamental difference in the kinds of things sensory imaginings and perceiving represent. I will present one now.

4. Imagining from the outside in

How is it that one and the same type of mental image can be used—appropriately—to sensorily imagine either an object (such as a suitcase) or an experience (such as perceptual experience as of a suitcase)? Peacocke's answer was that the same type of images is appropriate in each case because we are always imagining the same thing—a certain type of experience—even if we sometimes add, as a suppositional footnote, that other, non-mental things are present in the imagined scenario. I want to propose a different way of viewing things, by drawing on a framework for understanding the contents and correctness conditions of sensory (or "imagistic") imaginings that I've developed elsewhere (Langland-Hassan 2018, 2020, 2015).

Let us start with what it is that a mental image represents or "says" about the world. Suppose, for the sake of argument, that we can equally directly imagine experiences and ordinary objects, depending on our aims—that there is no asymmetry of the kind earlier discussed. Such a view would require that what an image represents is somehow different in different contexts. If that were the correct way of viewing things, then mental images could not be typed by their contents. For if they were so typed, the same type of image could not be used to (directly) imagine different things. This would entail that the most common means for typing mental states—viz., by the information they carry—is not open to us. Granted, there are other ways one might type images, such as by non-representational neurological or (for those so willing) phenomenological properties. However, I want to articulate a strategy for dealing with the puzzle of imagining experiences and imagining objects that allows for typing images by their content—by what the images are *about.* After all, to call something an image is already to presume that is an image *of* something. We should only conclude that, strictly speaking, images are *of nothing at all* if we can find no suitable way of typing them by their contents.

The key to typing by content, while also allowing that images directly represent non-mental entities, is to remove particularity from what they represent. Instead of thinking of the content of

a mental image of a suitcase as involving some particular suitcase, we can think of its content as akin to that of an indefinite description (i.e., a description beginning with "a" "an" or "some"), where an image's content may be something along the lines of *a suitcase of thus and such color and shape*. We can, for present purposes, leave open the precise properties that images are capable of representing in this way. It may be that a property such as being a suitcase is too abstract for this mode of representation and that images, strictly speaking, only represent superficial properties such as colors, shapes, and relative spatial relations (specifically, the types of properties that the human visual system evolved to discriminate and detect). In describing the properties images represent, I will default to such superficial properties, without any attached assumptions that they *must* be so limited. We can think of images, then, as representing types of outlays of perceptible properties—ways that shapes and colors may be laid out and interrelated in a three-dimensional scenario. A limitation I will put on the class of properties, however, is that they are not *mental* properties—and, in particular, they are not properties *of experiences*. Our visual systems did not evolve to detect or discriminate these, I will assume. Instead, they are the kinds of properties represented by perceptual experiences, just as the word 'red' may represent the property of being red without itself being red. This is all in keeping with the presumptive symmetry in the kinds of properties represented by imaginings and perceivings supported in Section 2. (And, while I will describe these contents using words of a natural language, we should not assume that the images themselves represent by means of language-like representations. The words are simply used as a way of gesturing at the kinds of properties images may represent.)

In what sense can we then still imagine particulars—such as the Arc de Triomphe? And what are we to make of our (apparent) ability to imagine experiences without also imaging objects, as when one imagines hallucinating a suitcase? My approach to these questions is the inverse to Peacocke's. Whenever we sensorily imagine, we always at least imagine some superficial non-mental properties, such as shapes and colors, through the use of mental imagery. This mental imagery combines with non-imagistic representations (akin to Peacocke's S-imaginings) to account for different particulars one might imagine—or even to assure that mental states, such as perceptual experiences, are imagined. However, on my framework, instead of seeing these non-imagistic "suppositions" as wholly distinct mental states from the imagery they accompany, we should view them as combining with the imagery into a single hybrid state, where the non-

imagistic component serves to pick out an object (or type of object) and the attached image serves to predicate properties to the object. *It is this total hybrid state that is the episode of sensory imagining proper.* Thus, the mental imagery used is just one component of the sensory imagining. (While the precise way in which images combine with non-imagistic representations in sensory imaginings is distinctive to my view, others have also proposed that the overall content of thoughts involving imagery involves contributions from non-imagistic states (Kaplan 1968; Fodor 1975; Tye 1991; Johnson-Laird 1983).)

Some examples will make this clearer. I will use **bold** to symbolize the component of the sensory imagining that is contributed by a mental image, while the remainder of the state's content—shown without bold—is represented by a non-imagistic state of some kind. Because the combined content is truth-evaluable, we can say that, together, these elements represent a proposition (even if, again, one of the representation's components is a mental image). So, what we might intuitively describe as "imagining the Arc de Triomphe" can be symbolized, in (1), as an imagining with the content:

(1) The Arc de Triomphe is **a large, grayish-white arch-shaped object.**

Here the first, non-imagistic component of the imagining serves to pick out a particular object, namely, the Arc de Triomphe. The imagistic component—whose content is only gestured at by the words in bold, and is not to be equated with their content—predicates the properties of being a large-grayish-white arch-shaped object to the Arc de Triomphe. Because the content of the image is akin to that of an indefinite description, the same type of image—typed by its content—can be used in other imaginings of this hybrid structure to imagine other particulars. Suppose, for instance, that there is a visually indistinguishable replica of the Arc de Triompe at Disney's Epcot Center. One might accurately imagine the Epcot arch with a sensory imagining of the form:

(2) The replica arch at Epcot Center is **a large, grayish-white arch-shaped object.**

This captures the legitimate sense in which the same type of image can be used to imagine different particulars. Granted, it leaves open the question of how the non-imagistic component of the state manages to refer to a particular. But this is a question that needs answering by all

sides, independent of any puzzles about sensory imagining, insofar as all sides are independently committed to the existence of non-imagistic thought about particulars.

Now to the key question of what it is to *directly* imagine something on this framework. If what we directly imagine is only the content of the image itself, then the view I am proposing entails that we never directly imagine particulars of any kind; we simply imagine types of three-dimensional outlays of superficial properties. Yet, we need not hold that what we directly imagine is simply the content of the image. On the view of sensory imagining I have proposed, the object of the imagining is accounted for by the non-imagistic portion of the content that accompanies the image, with the image playing a predicative role, attributing certain properties to the object or scene. (Recall that this view is independently supported by its ability to explain how we imagine different particulars with one and the same type of image.) Within this framework, we can hold that one directly imagines some *x* only if the properties represented by the image are represented as properties of *x.* For instance, in (1), the Arc de Triomphe is directly imagined because the properties represented by the image are represented as being properties of the Arc de Triomphe.

Importantly, this way of understanding direct imagining coheres with a natural view of direct visual perception. This is the view that we directly perceive some *x* when, thanks to being in appropriate causal contact with *x,* we represent superficial properties of *x.* For instance, I directly perceive the apple in front of me thanks to my visual experience's being caused by the apple's reflecting light to my retinas; and the apple is correctly perceived because the resulting experience attributes the properties of redness and apple-shaped-ness to its cause—namely, the apple. The only difference is that, in the case of sensory imagination, something other than appropriate casual contact with some *x* must determine the object of the imagining.

We can now consider the case of imagining hallucinating. It seems that I can imagine hallucinating the Arc de Triomphe. And this might seem to involve directly imagining an experience. After all, in the imagined scenario, there is no Arc de Triomphe. What is going on here? I suggest we understand the content of such an episode as:

(3) I am having an hallucinatory experience as of the Arc de Triomphe, **a large, grayish-white arch-shaped object.**

This imagining uses the same type of image—one that represents colors and shapes of an ordinary arch—as was used to imagine the Arc de Triomphe (and its replica at Epcot, in (2)). However, there is no arch in the imagined world. Instead, the full content of the episode entails that I am imagining a hallucinatory experience as of the Arc de Triomphe. Am I *directly* imagining the experience? No. The properties represented by the image are not represented as properties of the experience itself. After all, the hallucinatory experience is not being represented as a large, grayish-white and arch-shaped object. The image is here playing a very different role of specifying what the hallucinatory experience is *of* or *about* – of specifying properties of what it represents. We should instead say that I have *indirectly* imagined an experience, by means of forming an image of a large, grayish-white arch-shaped object.

This case of indirect imagining is an instance of a general phenomenon, where sensory images are used to represent properties of something other than the object of the imagining. There are many cases of indirect imagining that don't involve imagining experiences. When imagining that Joe has returned from his first visit to Paris and is still raving about the Arc de Triomphe, I might have a sensory imagining of the form:

(4) Joe won't stop talking about the Arc de Triomphe, **a large, grayish-white arch-shaped object.**

Here Joe's speech is indirectly sensorily imagined, insofar as a mental image is used in the imagining to characterize what his speech is *about*, without the image's attributing properties to the speech itself. This helps to highlight how indirect such forms of imagining really are. The image only tells us about Joe's speech to the extent that it lets us think about aspects of what he is talking about; it (obviously) does not represent any perceptible properties of the speech itself (even though the speech *has* aurally perceptible properties). Likewise, in the case of imagining a hallucinatory experience, the image only tells us about the nature of the experience to the extent that it tells us what the experience is about; it does not represent perceptible (or quasi-perceptible) color-like properties of the experience itself.

One might grant the possibility of indirectly imagining an experience in the way just described while insisting that it is also possible to directly imagine an experience. What would this

involve? On the present framework, it would need to be the case that the properties represented by the image are attributed to a mental state—a conscious experience—of some kind.

An example of this would take the form of:

(5) Joe's experience of the Arc de Triomphe was **a large, grayish-white arch-shaped object**

It is conceivable that someone would generate an imagining with this content. In fact, I think imaginings of this form are common. However, if generated in an assertoric mood, with the aim of accurately representing some entity (viz., Joe's experience), it would be misrepresentational. This is for the simple fact that Joe's experience is not a large, grayish-white, arch-shaped entity. The problem is, one might not recognize that they were engaged in this kind of misconceived imagining, however, because the imagery used does not represent its content via words of a natural language (such words are only used to gesture at the general types of contents had by these states, which are presumably not language-like representations). The fact that there is an evident conflict is thus possible to miss; it is possible to think, somewhat hazily, that the same imagery *can* represent properties of an experience *or* an object. This error is the root of the illusion that the conscious mind is somehow well-lit, or luminous. More on this in a moment.

First, there *is* a way to directly imagine experiences on the present framework, if we assume that experiences are neural processes and that we can directly imagine neural processes of that kind (e.g., by forming images that represent color and shape properties of conscious brain states). For instance, such an example might be:

(6) Joe's experience of the Arc de Triomphe was **a bit of pink, red, and tan neural issue.**

Of course, this is not the sort of thing people typically have in mind when they speak of "imagining experiences." Instead, they have mind something more like (5), where images of ordinary worldly scenes are used with the idea that they provide a characterization of what the experience itself is like. But such uses are inevitably *mis*representational when they represent the experience as having the properties that are the content of the image. They are veridical only if they characterize what the experience "is like" by telling us what the experience *is about*.

Now, in response to all this, one might say:

> So much the worse for your framework! It is a datum that we can directly imagine experiences. If your way of understanding the contents and representational contributions of mental imagery within sensory imaginings does not allow for this, we will need to use some other framework.

My first response in concessive: in articulating my approach an explaining some of its virtues, I have not shown that no other approach is possible that would allow us to directly imagine experiences. There are indeed views in the literature—which invoke imagery-involving, self-referential "phenomenal concepts"—that attempt to explain how mental imagery could do that special kind of work, and where such uses are essential to our knowledge of what it is like to have various experiences (Balog 2012; Block 2006; Papineau 2007). I have taken explicit aim at such views elsewhere (Langland-Hassan, 2018). Here I will limit myself to the positive project of explaining the advantages of my approach and (in the next section) to rebutting the view that it is strictly impossible to (directly) imagine an object *without also* imagining an experience.

First, we have seen that Peacocke is wrong that it is a "conceptual truth" that all sensory imagining requires "at least" imagining an experience. If this were a conceptual truth, it would not be possible to give a coherent account of imagining objects that does not involve the imagining of experiences. Whether one accepts the view I have sketched or not, it is not incoherent. Second, there is a way of treating the three cases Peacocke describes—imagining a suitcase, imagining a suitcase obscuring a cat, and imagining hallucinating a suitcase—that does not see the common factor as the imagining of a certain type of experience. Instead, the common factor is that each involves use of an image that represents a certain set of shapes and colors in a particular spatial configuration. The view I've offered requires a contribution from outside of the image itself to capture what is different in the three imaginings. Yet this sort of addition is also a feature of Peacocke's view (with its "S-imaginings").

In addition to providing an account of how particulars are imagined and of how the same type of image can be featured in imaginings of different particulars, I supported my approach through two analogies to perception. First, there is the empirical data reviewed in Section 2, which provides *prima facie* warrant for thinking that the representational capacities of mental imagery

and perceptual states will be of a similar kind (i.e., tracking similar superficial properties of the environment). Second, the account given of directly imagining an x mirrors the account one might give of directly perceiving an x: to directly imagine (or perceive) some *x* is to use a mental image (or perceptual experience) attribute to x superficial properties of a characteristic sort—including shape, color, and spatial-relational properties.

The last, and perhaps most powerful, reason to accept the sort of framework I have outlined is that it helps to dissolve the sense of mystery surrounding the mind-brain relation, in ways I shall now explain.

5. The dangerous desirability of light metaphors

We can now return to the puzzle with which we began: why it is that light metaphors seem *both* highly appropriate to the characterization of conscious experience *and* to capture the element of consciousness that eludes the grasp of ordinary mechanistic and scientific explanation? On the framework I have recommended, when we (try to) directly imagine experiences, we engage in *misrepresentational* imaginings of the kind in (5), where ordinary color and shape properties are (unwittingly) attributed to mental states. The tendency to do so has two key effects: first, because such imagery really does represent colors and shapes—and, in general, well-lit environments—it will seem to us that there is something "illuminated" about the nature of the experiences so imagined. The lights will indeed seem to be *on* within the mind. Second, because the neural processes that are, in fact, our conscious experiences completely lack the properties we are (again, unwittingly) attributing to them, it will seem mysterious how they could have them. The relationship between the kinds of brain processes studied in neuroscience and the properties we attribute with such imagery will seem entirely contingent; it will be very puzzling indeed how any of the physical properties of the brain could just *be* properties of the kind attributed by the mental imagery.[6] The reason is not that property dualism is true, but that

[6] This diagnosis of the sense of contingency we may have concerning the mind/brain relation was initially inspired by remarks of Thomas Nagel in footnote 11 to his (1974) "What is it like to be a bat?" There he distinguishes two forms of imagination—"sympathetic" and "perceptual" imagining—and proposes that differences in the way they operate would account for the appearance of contingency in the mind/brain relation, even if the relation were one of identity (see also Hill (1997)). A crucial difference in my and Nagel's accounts is that Nagel does not find an *error* in any of the relevant imaginings. They are simply represented as two disparate ways of knowing things. This leaves the core puzzle about consciousness intact—which is unsurprising, given the main thrust of the famous paper

our minds entirely lack the shape, color, and spatial-relational properties attributed by imagery whenever we (think we) "directly" imagine experiences. In short, the hard problem of consciousness is, at least in part[7], the result of an illusion generated by inappropriate efforts to sensorily imagine experiences.

As this is an error-theory of a kind—one which finds a pervasive mistake in the way people are inclined to imagine experiences—I owe an account of why we would make such an error and of why we might not recognize that we are doing so. In response to the first question: in most contexts, when we think about the thoughts of others (or even our own thoughts) we are mainly interested in what those thoughts (or "experiences") are *of* or *about*. We are not ordinarily interested in the intrinsic, neural (or "vehicular") properties of the thoughts. This why, in ordinary life, when we are asked to describe our thoughts, we immediately speak of what our thoughts are *about.* This primary interest in the objects of others' thoughts leads us to closely associate images of whatever someone is thinking about with the person's thoughts themselves. For instance, when wondering what Sally sees as she peers into a covered box—and suspecting it is a turtle—we may form an image of a turtle and, in a sense, tag it to Sally's mind, thinking:

(7) Sally's perceptual experience is like this: **a little green and yellow turtle-shaped thing.**

Of course, Sally's experience is not *itself* a green and yellow turtle-shaped thing. It does not have the superficial properties of a turtle. But the fact that it doesn't does not disrupt the purposes to which we put the image. Our goal is simply to know what Sally is perceiving and to keep track of that fact. The imagining in (7) allows us to do so. Supposing there really is a green and yellow turtle in the box, whether (7) is misrepresentational will depend on what

in which this footnote appears. For more on the relation between Nagel's footnote and my view of the objects of sensory imagination, see Langland-Hassan (2018).

[7] I do not think that the every aspect of the mind-body problem is due to the kind of mistaken imaginings I've been describing here, even if they are an important part of the equation. Also relevant are epistemological phenomena such as the (apparently) privileged access we have to our own minds, and the (apparently) greater certainty we can have of our minds than we can have about other things in the world. These phenomena require separate treatment (as in, e.g. Byrne (2018)). Also relevant in this context are intuitions about the nature of the self and personal identity, as explored, e.g., in Parfit (1984).

precisely is meant by "like this." If "like this" means "has the following properties," then (7) misrepresentational. However, if "like this" is more accurately disambiguated as "is of" then (7) is veridical. In the midst of everyday life, we do not bother with this disambiguation. The thought serves its purpose in associating the person with the content of the image. The fact that thoughts like (7) can (depending on their disambiguation) involve a predicative error only generates confusion when one turns to the philosophical and scientific project of understanding the relationship between neural processes and mental states. By that time, the heuristic of tracking others' thoughts by associating mental images of their thought contents with their thoughts themselves has become so frequently and fruitfully employed that it has an aura of indubitability. Thus, while the error is fairly easy to describe, convincing others that it really *is* an error remains an uphill climb.

Fortunately, there is support from developmental psychology for this general account of why and how we come to misimagine experiences. In three-year-olds, the use of a "thought bubble" heuristic has been shown to dramatically improve performance on traditional theory of mind tasks (Wellman, Hollander, and Schult 1996; Kerr and Durkin 2004). In these studies, the experimenters use thought bubbles above the heads of depicted or puppet characters to explain the idea that different people can have different (sometimes conflicting) thoughts. When prompted with this imagistic way of "keeping track" of what others are thinking about, the ability of these young children to correctly ascribe false beliefs to characters in the stories increases significantly. The experimenters can be seen as jump-starting the use of a heuristic that is eventually mastered by older children and adults in the ordinary course of development. Yet it is a heuristic that can very easily lead to precisely the sort of error I have described. In a similarly supporting line of research, it has been shown that children with Autism Spectrum Disorder (ASD)—a condition characterized by difficulties in social cognition—markedly improve on experimental mindreading tasks when trained using a "pictures in the head" paradigm, where they are taught to associate pictures of what a person is thinking about with the person's head (Wellman et al. 2002; McGregor, Whiten, and Blackburn 1998; Parsons and Mitchell 1999). Payntner & Peterson (2013) showed that performance in this population remained improved even three weeks after such training.

These studies are a nice validation of the idea that associating imagery with someone's thoughts—and, indeed, with someone's *head*—is a powerful and commonly used heuristic for understanding and tracking the contents of other minds. The harmless and helpful heuristic only turns pernicious when the thought that someone's experience is "like this" comes to be seen as a comment on the properties of the experience itself, and not the object of the experience. The mind then appears oddly luminous.

6. Must we imagine experiences? Against the Dependency Thesis.

I have provided a deflationary diagnosis of light metaphors in the description of consciousness—one that sees their popularity as grounded in a common way of misimagining experiences. If correct, it has the advantage of explaining why such metaphors are attractive, while dissipating some of the mystery surrounding the mind/brain relation. The strategy outlined depends on a particular conception of the structure of sensory imaginings, according to which experiences are only derivatively imaginable, and according to which one can imagine objects without imagining any experiences at all. This approach clashes with views that hold both experiences and ordinary objects to be directly imaginable (which, as earlier noted, I address, elsewhere (Langland-Hassan, 2018), and with views which hold that the imagining of objects *always* depends upon the imaging of an experience. The latter view has come to be known as the Dependency Thesis, according to which imagining some *x* always requires imagining an experience of *x.* We have already discussed and rebutted Peacocke's (1985) defense of the Dependency Thesis. In the last section of this chapter, I want to address a more developed version of that sort of view put forward by M.G.F. Martin (2002). Defusing Martin's argument for the Dependency Thesis is crucial for the coherence of framework I have defended.[8]

Martin (2002) argues that when we imagine ordinary objects (e.g., an apple) we can only do so "through imagining a visual experience as of an apple" (2002, p. 411). He takes visualizing as the central form of sensory imagination. "In visualizing a tree," he proposes, "one is imagining visually experiencing a tree" (2002, p. 411). Put another way—and introducing the notion of a

[8] For important critiques of Martin's Dependency Thesis that are independent of the points raised here, see Noordhof (2002, 2018).

"point of view"—he says, "If we are to get right how we have visualized things, then we need to introduce the point of view and the experience into the imagined situation" (*ibid*.).

Now, because Martin, like most contemporary philosophers, does not think that ordinary perceptual experiences represent mental entities, his argument that sensory imagination requires imagining experiences must appeal to some *difference* between perception and imagination. And, by the same token, whatever feature sensory imagination has in virtue of which it involves representing experiences must not be a feature that it shares with perceptual experience.

For instance, the fact that objects are visually perceived from a point of view (whatever this entails) does not for Martin require that visual experiences are seen or represented in perception. But, he argues, matters are different in the case of imagination. His central argument involves imagining "a world which contains merely two spots of light" (p. 408). He asks us to imagine a red spot of light being to the left, and a green spot of light to the right. He then claims that we can also visualize the reverse—the green spot of light to the left, and the red spot of light to the right. However, under the supposition that these spots of light are the only things in the imagined world, there "can be no difference between the two situations." The two situations "count as different only where there is a point of view relative to which the one object is to the left and the other to the right, or *vice versa*." That is, "if we absent a point of view from the imagined scene, then what appears in visualizing to be a difference in the scene imagined, and not just a difference in one's state of mind cannot be so" (p. 409). Thus, respecting the phenomenological datum that something different has been imagined in the two cases requires us "to imagine a point of view within the scene," and, therefore, one "must be imagining an experience within the scene" (p. 409).

One might object that the perspectival nature of imagination Martin appeals to holds for perceptual experience as well and that he therefore runs afoul of the self-imposed constraint not to conclude that perceptual experience requires representing something mental. While he admits that perceptual experience has similar perspectival elements—representing things "to the left" and "to the right" and so on—matters are different in the case of sensory imagination, he thinks, because "in imagining something as to the left one does not thereby imagine [it] as in one's actual environment on the left." Rather, the perspectival aspect of one's visualizing relates "not

to one's *actual* situation, but rather to the *imagined* situation." (p. 410). "The red light is imagined as before and to the left of the point of view within the imagined situation by being imagined as presented to a point of view within that situation, and hence as being experienced as to the left from that point of view." In this way, "an experience-relative aspect of a visualized scene, how it and its elements are oriented, is imagined through imagining an experience with the appropriate property."

I have quoted Martin at length, so as not to misrepresent his subtle argument. We are trying to determine how sensory imagining is *different* from perceptual experience such that only it involves representing an experience. Martin's idea is that in perception the actual environment serves as an independent frame within which an object is represented as being in a certain direction "to the left" or "to the right". Thus, a "point of view"—an experience *to which* the objects are so related—need not also be represented. The objects are to the left or right of one's body *in the actual environment*; no need to mention an experience. But in the case of imagining two spots of light in the two different ways, they are not represented as being in one's actual environment to the left and right. And, supposing they are the only things in their imagined world, they are not to the left or right of some other (non-mental) object. So, in imagining them in these two ways, we must be imagining not just the lights themselves but the lights *as experienced* in two different ways.

To assess this argument, we can begin by asking why Martin has chosen something so odd as a "spot of light" in his example. What, exactly, is a "mere" spot of light? Is it a bit of light that is neither projected from nor reflected by something? Is it just a collection of photons hanging around in the shape of a circle…not moving anywhere *at the speed of light*? One might question the conceptual coherence and metaphysical possibility of such an entity. I will return to this.

Backing up a moment, Martin begins his argument less controversially, appealing to the imagining of two actual, everyday lights (they may as well be traffic lights). We imagine one to the left, one to the right, and then vice versa. Again there is a difference in the two scenes imagined. But if what we are imagining are real, three-dimensional traffic lights, there will always be differences in the scene itself that can plausibly explain the phenomenal difference in the two imaginings, without need to appeal to an experience that is imagined in the scene. Either

a different part of at least one light will be imagined in the two cases (e.g. their back sides are now represented), or the parts that are imagined will be imagined as having different spatial relations to each other.

To see this, imagine two uniformly luminous discs, one green, the other red. Now suppose the discs switch sides without either turning clockwise or counter-clockwise or flipping to reveal its back side—the green just hops to the alternate side of the red such that the part of the green disc that had been furthest from the red disc is now the side of the green disc closest to the red disc. This is a case of imagining (what we could otherwise describe as) a left/right switch of the discs, where there is a clear difference in the scene imagined. Notice that, to capture that difference, there is no need to mention an experience, despite there being nothing in the imagined world other than the lights. In fact, to capture the difference, there is no need to mention "left" or "rightness" in general. We just imagine them situated in one way, and then the other, where the difference in the scenes consists in differences in the spatial relationships of parts of the two discs to each other.

To consider another case, suppose that the same "hop" occurs, where the green disc hops to the other side of the red one. But suppose also that the green disc spins 180 degrees on its axis such that the part of the green light that was originally closest the red light is again closest to the red light. Now the spatial relations of parts of one disc to parts of the other have been duplicated in the two scenes[9] and there has been (what we could otherwise describe as) a left/right switch. But there is still an obvious non-experiential difference in what is imagined: you are imagining a different *side* of the green disc than you were before.

The lesson: so long as we are imagining three-dimensional objects, there will always be differences in the scene itself that account for the sort of "phenomenal" difference Martin adduces, without our needing to appeal to an imagined experience—or even to left-ness or right-ness. Martin comes close to recognizing this in a footnote, explaining that his argument must be taken to pertain to *mere* spots of light (and not *lights*) because:

[9] Well, not really, for we are overlooking the *thickness* of each disc and *its* implications

> If we think instead of objects which have internal orientation (i.e. which themselves have a top and bottom or front and back and hence a left or right), then differences between the two situations will turn up from whether the red object's left side is adjacent to the green object's right side, or *vice versa* (fn. 36, p. 409).

Yet the point pertains not only to objects that have an "internal orientation," (i.e. anything with an obvious front and back) but to objects *generally*. Nothing in the example I provided assumed a top, bottom, left, right, front, or back to the lights themselves. Rather, being physical objects, the lights have *parts*. And what we think of, intuitively, as the lights switching sides—green going from left of red to its right—is always accompanied either by changes in the distance among various parts of the lights to each other, or by changes in which parts of the lights are being imagined. So, if the spots of light have any parts, and if they extend into three spatial dimensions, there will be differences in the two imagined situations after all, without appeal to an experience. If, on the other hand, spots of light have no parts and are (merely hypothetical) two-dimensional entities, then they are not the kinds of things we could ever perceive. The example would then show, at best, that when it comes to imagining things that are imperceptible, we must imagine an experience. Whether or not this is coherent, it is a far cry from Martin's stronger conclusion that all sensory imaginings require the imagining of an experience. Thus, Martin has not given us reason to doubt that ordinary objects are directly imaginable in the way described last section. The deflationary account of imagining experiences remains intact.

7. Blink's return

Why do light metaphors track intuitions about consciousness? I have argued that people take themselves to directly imagine experiences and that such imaginings involve mental imagery. As imagery invariably represents well-lit objects, this naturally leads people to think that there is something *almost* literally well-lit about the mind. This leads in turn to the conviction that brain states have certain kinds of "phenomenal" properties that they cannot possibly have, and thus to mysterianism about the mind-brain relation. I argued for this illusionist view by motivating a presumption in favor of (visual) sensory imagination's representing in the manner of visual perception and then providing a framework for understanding the contents of sensory imaginings—one that shows both how the same type

of image can be used to imagine different objects and why ordinary objects, but not experiences, are directly imaginable. This account was bolstered by an explanation of why we are inclined to fall into the error of thinking we can directly imagine experiences, which appealed to certain ways imagery is reflexively used to keep track of other minds. Finally, I rebutted M.G.F. Martin's argument that the imagining of ordinary objects depends upon the imagining of experiences.

I close with a lingering question. Recall Blink, the 21st Century robot (capable of color and shape discrimination, lively conversation, GPS-guided travel, and more) introduced at this chapter's outset. For Blink and other such robots, *all is dark* inside. At least, that much is conventional wisdom. But if my proposal about light metaphors is correct, why does Blink's cognition seem so dim? Why don't we naturally use mental imagery to think about *Blink's* "thinking" too, and so conclude that Blink's thoughts are equally phosphorescent, equally luminous? Why the sudden chauvinism with respect to Blink—or have I completely misdiagnosed the relation between mental imagery and light metaphors?

In response, I don't think that our approach to Blink (to the extent that this *is* our approach) is at all straightforward. If there is a chauvinism—if Blink alone is in the dark—the chauvinism is unwarranted. But pinpointing the exact reasons and motivations for it is not straightforward. So, I end with a not straightforward reply.

Last we heard of Blink, he was headed to California. No doubt you've been wondering how he made out—whether he met gentle people there, what it was like *for* him to do so. I'm happy to report that he is getting on very well. He has joined you and your companion at the beach for a long day of fun: catching waves, playing paddleball, discriminating colors and shapes. The air is cooling. A magnificent sunset unfolds. Far off, a jet crosses the sky, leaving behind it a bright trail of orange smoke, as though turning back the peel of a ripe nectarine. Richard Branson appears above you in a giant hot air balloon, then floats away again, continuing on in his remarkable quest.

You look at Blink and think about all that he is "seeing": the indelible hues, the receding hot air balloon, the rolling tide. You look to your companion, who is similarly taken with the magical scene. Considering her experience, you think: all of *that*—the sea, the sky, the

gliding jet, the Virgin mogul—all of that is illuminated in *there,* in her mind, a technicolor movie on a boundless, transparent screen…just as it is in mine. The richness of human experience!

And then you look again at Blink who, truth be told, appears less rapt. "Orange sky," comments Blink. "Jet," he says. "Hot air balloon in the way," he notes. Having still the rich visual image of all that is before you in mind, lingering from your meditation on the visual experience of your companion, you wonder: should I attribute *all this* to Blink? All of this richness, this luminosity? Or should I deny him this bountiful image since, as they say, all is dark within him?

As you mull this over, Blink turns to look at you. Blink blinks. "If it's dark in me," he says, "it's dark in you."

References

Aydede, Murat. 2019. Pain. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford University.

Balog, Katalin. 2012. "Acquaintance and the mind-body problem." In *New Perspectives on Type Identity: The Mental and the Physical*, edited by Simone Gozzano and Christopher S. Hill, 16-42. Cambridge University Press.

Bartolomeo, P. 2008. "The neural correlates of visual mental imagery: an ongoing debate." *Cortex* 44:107-108.

Block, Ned. 1995. "On a confusion about a function of consciousness." *Brain and Behavioral Sciences* 18 (2):227–247.

Block, Ned. 2006. "Max Black's objection to mind-body identity." *Oxford Review of Metaphysics* 3:3-78.

Brooks, L. 1967. "The suppression of visualization by reading." *Journal of Experimental Psychology* 19:289-299.

Byrne, Alex. 2001. "Intentionalism defended." *The Philosophical Review* 110 (2):199-240.

Byrne, Alex. 2018. *Transparency and Self-Knowledge*. Oxford: Oxford University Press.

Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford Universtiy Press.

Craver-Lemley, Catherine, and Adam Reeves. 1992. "How visual imagery interferes with vision." *Psychological Review* 99 (4):633-649.

Damasio, Antonio. 1999. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Orlando, FL: Harcourt.

Dretske, Fred. 2003. "How do you know you are not a zombie?" In *Privileged Access: Philosophical Accounts of Self-Knowledge*, edited by Brie Gertler, 1--14. Ashgate.

Fodor, Jerry A. 1975. *The Language of Thought*. New York: Crowell.

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

Ganis, G., W.L. Thompson, and S. Kosslyn. 2004. "Brain areas underlying visual mental imagery and visual perception: an fMRI study." *Cognitive Brain Research* 20:226-241.

Hill, Christopher. 1997. "Imaginability, conceivability, possibility and the mind-body problem." *Philosophical Studies* 87:61-85.

Johnson-Laird, P.N. 1983. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press.

Kaplan, David. 1968. "Quantifying in." *Synthese* 19 (1-2):178-214.

Kerr, S., and K. Durkin. 2004. "Understanding of thought bubbles as mental representations in children with autism: implications for theory of mind." *J Autism Dev Disord* 34 (6):637-48.

Koch, Christof. 2004. *The Quest for Consciousness*. Englewood, CO: Roberts & Company.

Kosslyn, S., A. Pascual-Leone, O. Felician, S. Camposano, J.P. Keenan, W.L. Thompson, G. Ganis, K.E. Sukel, and N.M. Alpert. 1999. "The role of area 17 in visual imagery: Convergent evidence from PET and rTMS." *Science* 2 (5411):167-170.

Langland-Hassan, Peter. 2015. "Imaginative Attitudes." *Philosophy and Phenomenological Research* 90 (3):664-686. doi: 10.1111/phpr.12115.

Langland-Hassan, Peter. 2017. "Pain and Incorrigibility." In *The Routledge Handbook of the Philosophy of Pain*, edited by J. Corns, 245-254. London: Routledge.

Langland-Hassan, Peter. 2018. "Imagining Experiences." *Noûs* 52 (3):561-586. doi: doi:10.1111/nous.12167.

Langland-Hassan, Peter. 2020. *Explaining Imagination*. Oxford: Oxford University Press.

Levine, D.N., J. Warach, and M.J. Farah. 1985. "Two visual systems in mental imagery: Dissociation of "what" and "where" in imagery disorders due to bilateral posterior cerebral lesions." *Neurology* 35:1010-1018.
Martin, M.G.F. 2002. "The Transparency of Experience." *Mind and Language* 17 (4):376-425.
McGinn, Colin. 2000. *The Mysterious Flame: Conscious Minds in a Material World*. New York, NY: Basic books.
McGregor, Evelyn, Andrew Whiten, and Pennie Blackburn. 1998. "Transfer of the Picture-in-the-Head Analogy to Natural Contexts to Aid False Belief Understanding in Autism." *Autism* 2 (4):367-387. doi: 10.1177/1362361398024004.
Mechelli, A., C.J. Price, K.J. Friston, and A. Ishai. 2004. "Where bottom-up meets top-down: neuronal interactions during perception and imagery." *Cerbral Cortex* 14:1256-1265.
Nagel, Thomas. 1974. "What is it like to be a bat?" *The Philosophical Review* 83 (4):435-450.
Noordhof, Paul. 2002. "Imagining objects and imagining experiences." *Mind and Language* 17 (4):426-455.
Noordhof, Paul. 2018. "Imaginative Content." In *Perceptual Imagination and Perceptual Memory*, edited by F. Macpherson and F. Dorsch, 96-129. Oxford: Oxford University Press.
O'Craven, K.M., and N. Kanwisher. 2000. "Mental imagery of faces and places activates corresponding stimulus-specific brain regions." *Journal of Cognitive Neuroscience* 12:1013-1023.
Papineau, D. 2007. "Phenomenal and perceptual concepts." In *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, edited by Torin Alter and Sven Walter, 111--144. Oxford University Press.
Parfit, Derek. 1984. *Reasons and persons*. Oxford Oxfordshire: Clarendon Press.
Parsons, Sarah, and Peter Mitchell. 1999. "What Children with Autism Understand about Thoughts and Thought Bubbles." *Autism* 3 (1):17-38. doi: 10.1177/1362361399003001003.
Paynter, Jessica, and Candida C. Peterson. 2013. "Further evidence of benefits of thought-bubble training for theory of mind development in children with autism spectrum disorders." *Research in Autism Spectrum Disorders* 7 (2):344-348. doi: http://dx.doi.org/10.1016/j.rasd.2012.10.001.
Rosenthal, D.M. 2004. "Varieties of higher-order theory." In *Higher-order theories of consciousness: an anthology*, edited by Rocco Gennaro, 17-44. Amsterdam: John Benjamins.
Ryle, Gilbert. 1949/2009. *The Concept of Mind*. New York: Routledge.
Sirigu, A., L. Cohen, J.R. Duhamel, B. Pillon, B. Dubois, Y. Agid, and C. Pierrot-Deseilligny. 1995. "Congruent unilateral impairments for real and imagined hand movements." *NeuroReport* 6:997-1001.
Stokes, M., R. Thompson, R. Cusack, and J. Duncan. 2009. "Top-down activation of shape-specific population codes in visual cortex during mental imagery." *The Journal of Neuroscience* 29 (4):1565-1572.
Tye, Michael. 1991. *The imagery debate*, *Representation and mind*. Cambridge, Mass.: MIT Press.
Wellman, H. M., S. Baron-Cohen, R. Caswell, J. C. Gomez, J. Swettenham, E. Toye, and K. Lagattuta. 2002. "Thought-bubbles help children with autism acquire an alternative to a theory of mind." *Autism* 6 (4):343-63.
Wellman, H. M., M. Hollander, and C. A. Schult. 1996. "Young children's understanding of thought bubbles and of thoughts." *Child Dev* 67 (3):768-88.