

Mundane hallucinations and new wave relationalism

Jacob Beck

Department of Philosophy & Centre for
Vision Research

Correspondence

Jacob Beck, York University Toronto,
Canada.

Email: jbeck@yorku.ca

Abstract

Relationalism maintains that mind-independent objects are essential constituents of veridical perceptual experiences. According to the argument from hallucination, relationalism is undermined by perfect hallucinations, experiences that are introspectively indistinguishable from veridical perceptual experiences but lack an object. Recently, a new wave of relationalists have responded by questioning whether perfect hallucinations are possible: what seem to be perfect hallucinations may really be something else, such as illusions, veridical experiences of non-obvious objects, or experiences that are not genuinely possible. This paper argues that however well new wave relationalism may handle brains in vats, drug users “seeing” pink elephants, and other extraordinary hallucinations, it struggles to accommodate mundane hallucinations, such as “hearing” your child cry out from the room down the hall when she is actually sound asleep or “feeling” vibrations on your thigh even when your phone isn’t in your pocket. Mundane hallucinations are best explained as byproducts of noise in the perceptual system, and noise-induced hallucinations are resistant to the strategies that new wave relationalists deploy to explain away other hallucinations. Mundane hallucinations can thus underpin an especially powerful version of the argument from hallucination.

When philosophers discuss hallucinations, the examples they conjure are often extraordinary: a brain in a vat that is stimulated by a supercomputer to have human-like experiences; a disembodied “space soul” floating through the cosmos with experiences that are qualitatively identical to your own at this very moment; a drug user who consumes a hallucinogen and “sees” a pink elephant in her living room. But most actual hallucinations are mundane: “hearing” your child cry out from the room down the hall when he is really sound asleep; “feeling” phantom phone vibrations on your thigh even when your pocket is empty; “seeing” non-existent shooting stars in your periphery when you’re searching the night sky for real ones. For some philosophical projects, the focus on extraordinary hallucinations serves a useful function. For others, it threatens to cause trouble. This article is about one manifestation of that trouble.

1 | NEW WAVE RELATIONALISM

Let us define *relationalism* as the view that mind-independent objects are essential constituents of veridical perceptual experience, and a *perfect hallucination* as an experience that is introspectively indistinguishable from a veridical perceptual experience but lacks a mind-independent object. The *argument from hallucination* purports to make trouble for relationalism through two claims: that perfect hallucinations are possible; and that perfect hallucinations and veridical experiences have the same nature. Since mind-independent objects are not essential constituents of perfect hallucinations, it follows that they cannot be essential constituents of veridical perceptual experiences either. Relationalism is imperiled.

Most relationalists respond by embracing *disjunctivism*, which rejects the claim of a common nature. Disjunctivists maintain that perfect hallucinations and veridical perceptual experiences differ in kind even though they are indiscriminable from one another (Brewer, 2011; Byrne & Logue, 2008; Campbell, 2002; Fish, 2009; Hinton, 1967; Martin, 2002, 2004, 2006; McDowell, 1982; Snowden, 1980; Soteriou, 2016). But disjunctivism has been subjected to powerful critiques for, among other things, providing an inadequate account of hallucination’s phenomenal character (Hellie, 2007; Masrour, 2020; Pautz, 2010, 2011; Schellenberg, 2010, 2011, 2018; Siegel, 2004, 2008; Sturgeon, 1998) and being in tension with the science of perception (Burge, 2005, 2010; Pautz, 2017, 2021).

Recently, a new wave of relationalists have thus grown disenchanted with disjunctivism and begun to explore an alternative approach to the argument from hallucination: denying that perfect hallucinations are possible. As many of these new wave relationalists see it, disjunctivists have been tying themselves into knots to accommodate a fantasy. Any experience that lacked an object could not possibly be indiscriminable from a veridical perceptual experience. To explain away alleged perfect hallucinations, philosophers such as Alva Noë (2004; 2007), Thomas Raleigh (2014), Rami Ali (2018), Farid Masrour (2020), Michael Barkasi (2020), and Alex Byrne and Riccardo Manzotti (draft), deploy two strategies. The *denial strategy* denies that putative perfect hallucinations really are indiscriminable from veridical perceptual experiences. Thus, it is claimed that dreams have an ethereal quality that makes them in-principle discriminable from veridical perceptual experiences. The hallucinations of a space soul are also subjected to the denial strategy on the grounds that our intuitions about the possibility of such outlandish scenarios shouldn’t be trusted. For other hallucinations, new wave relationalists adopt the *redescription strategy*: they redescribe hallucinations as experiences that have a mind-independent object. Thus, according to some new wave relationalists, the brain in a vat doesn’t hallucinate a room with a desk; it veridically perceives some lines of code in the supercomputer that is stimulating it. Similarly, the drug

user doesn't hallucinate an elephant in his room; he suffers a visual illusion in which he misattributes the property *containing an elephant* to his (otherwise veridically perceived) room.

In this way, new wave relationalists seek to divide and conquer (see esp. Masrour, 2020). Some hallucinations are *not* indiscriminable from veridical perceptual experiences (the denial strategy). The rest may be indiscriminable from veridical perceptual experiences, but can be redescribed as experiences with non-obvious objects (the redescription strategy). For those who are attracted to relationalism, the appeal of new wave relationalism is easy to see. It promises to avoid disjunctivism's gymnastic account of perfect hallucinations while retaining its explanations of how perception yields direct epistemic (McDowell, 1982), semantic (Campbell, 2002), and phenomenal (Martin, 2002; Fish, 2009; Logue, 2012) access to the external world.

In defending their view, new wave relationalists have disproportionately focused on extraordinary hallucinations. Thus, while they puzzle over brains in vats, space souls, and drug users who experience pink elephants, they never consider phantom phone vibrations, "heard" voices, or other mundane hallucinations. This selective emphasis matters because mundane hallucinations pose a special challenge. As I will argue in Sections 3–4, they are readily handled by neither the denial strategy nor the redescription strategy. On the one hand, they are not like dreams or the hallucinations of space souls because there is no plausible way to deny that they exist and are indiscriminable from veridical perceptual experiences. On the other hand, they are not readily redescribed because there is nothing that can plausibly serve as their objects. Mundane hallucinations are thus uniquely troubling because they thread the needle between new wave relationalism's two main strategies for accommodating alleged hallucinations. For this and other reasons, they can be used to construct an especially powerful version of the argument from hallucination, as I will show in Section 5.

While I will argue that mundane hallucinations are not easily accommodated by new wave relationalism, I will not claim that they are impossible to accommodate. There are moves available to the new wave relationalist. But as I will show, these moves involve commitments to further theses that are, at the very least, highly controversial. The arguments to come should thus be viewed as a challenge. If new wave relationalists want to deny the possibility of perfect hallucinations, they must explain how they will handle mundane hallucinations—an explanation that, up until now, they have not furnished, and that, however achieved, is likely to incur significant costs.

If mundane hallucinations cause such trouble for new wave relationalism, it might seem that relationalists should retreat to disjunctivism. This is, I suspect, their best available option. If so, then the main contribution of this paper is to box relationalists into disjunctivism. For those who are sympathetic to disjunctivism, this won't be an unwelcome result. But for those who are, like me, impressed by critiques of disjunctivism, it may instead inspire the rejection of relationalism in any form.

2 | THE SCIENCE OF MUNDANE HALLUCINATIONS

One reason that mundane hallucinations are so much harder for new wave relationalists to accommodate than extraordinary hallucinations is that perception science yields deep insights into how they are produced, and many of the moves that new wave relationalists make to explain away extraordinary hallucinations do not comport with what these insights teach us about mundane hallucinations. In particular, I will argue that mundane hallucinations are best explained as byproducts of noise in the perceptual system. To understand these hallucinations, it is thus necessary to understand perceptual noise.




FIGURE 1 (a) If the same stimulus intensity always elicited the same response, for any given intensity value a perceiver would either detect it or not, leading to a step-function. (b) In reality, detection functions are S-shaped

2.1 | Perceptual Noise

What is the quietest sound that you can hear? A simple model for thinking about this question posits two variables: a stimulus intensity, I , and a sensitivity threshold, T . When $I > T$, you perceptually experience the stimulus. Otherwise, you fail to experience the stimulus. Thus, according to this simple model, the quietest sound you can hear is a sound whose intensity just surpasses the sensitivity threshold of your auditory system.

There are various problems with this simple model, but one of the main difficulties is that it fails to account for noise, random fluctuations that are unrelated to the stimulus. In perceptual psychology, it is widely assumed that noise is an ineliminable aspect of perceptual processing. One consideration that motivates this assumption derives from response patterns in threshold detection tasks. Responses to questions such as “Did you hear a tone?” fluctuate variably even when the signal and all other controllable background conditions remain constant. For example, if we try to determine the quietest sound you can hear by playing a series of tones of variable intensity until you identify the minimal intensity tone that is still audible, we’ll find that the same tone will not elicit the same response every time it is played. Rather, your judgments of whether you heard a given tone will be probabilistic. They’ll look like Figure 1b rather than Figure 1a. On the assumption that your experiences are determined not only by the strength of the stimulus itself, but also by noise, this finding makes perfect sense.

Another reason to believe in the existence of perceptual noise is that the best models of perceivers’ response behavior in psychophysical detection tasks explicitly assume that signals are always corrupted by noise. For example, according to signal detection theory (Green & Swets, 1966), responses are a function of two variables: d' and β . d' is a measure of detectability—that is, the extent to which the signal-plus-noise exceeds the noise alone. So the greater the value of d' , the easier it is to detect the signal. β corresponds to the perceiver’s “bias,” or criterion for responding that a signal is present. As β increases, the perceiver becomes more conservative—that is, the signal-plus-noise has to reach a higher value before the perceiver is willing to respond that a signal is present. For our purposes, the details don’t really matter beyond the fact that signal detection theory presupposes the existence of noise. It is hard to overstate the centrality of signal detection theory to contemporary perceptual psychology. It “has become the modern standard for studying a wide variety of perceptual and memory phenomena” (Palmer, 1999, p. 668) and “may well be regarded as the most towering achievement of basic psychological research of the last half

century” (Estes 2002, p. 15). The fact that it posits perceptual noise thus provides further reason to believe that such noise really exists.¹

While signal detection theory posits noise, strictly speaking it is silent on whether the source of that noise is external or internal to the mind. But it is widely assumed that both types of noise are ubiquitous. For example, when you’re in a tone detection task, the room might not be perfectly silent apart from the tone. At the very least, your breathing will be faintly audible (a source of external noise). Additionally, internal processing is grounded in our biology, which admits of random fluctuations—in blood flow, temperature, metabolic activity, membrane potentials, and neural firing rates, for example. The visual cortex contains neurons that are tuned to specific orientations, such as “vertical neurons” that fire at their maximal rate when a line that is vertically oriented falls within its receptive field, and at decreasing rates as the line’s orientation moves away from verticality (Hubel & Wiesel, 1962). But a vertical neuron doesn’t fire at *exactly* the same rate every time a vertical line (say) is presented in its receptive field. Rather, the neuron’s firing rate in response to a vertical line varies around a mean value. The precise rate at which the vertical neuron fires in response to a vertical line is thus noisy (Tomko & Crapper, 1974). The sources of this neural noise are many and diverse, but include the stochastic release of neurotransmitter vesicles, variability in synaptic conductance, and “channel noise,” the random opening and closing of ion channels (Faisal et al., 2008; Rusakov et al., 2020). These biochemical and electrochemical noise sources are sub-cellular; but they “can significantly alter whole-cell responses,” such as the generation of action potentials (Faisal et al., 2008, p. 294).

Another example of internal noise that has the advantage of being fairly well understood concerns the earliest stages of light-sensitive receptors in the retina. While retinal photoreceptors are activated by light, they also exhibit spontaneous activation in the dark. This spontaneous activation results from the inherent instability of rhodopsin, the protein molecule responsible for visual phototransduction (Ashmore & Falk, 1977; Baylor et al., 1980). While rhodopsin reliably isomerizes (reorganizes into a new molecule with the same atoms but in a different arrangement) in response to light, it also occasionally isomerizes spontaneously. As a result, “Retinal photoreceptors are noisy. They generate discrete electrical events in the dark indistinguishable from those evoked by light” (Barlow et al., 1993, p. 64).

Perception science is thus committed to the existence of internal noise. But can such noise in individual receptors and neurons make a difference to behavior and experience? One reason to think so is that minor interventions can induce a noticeable change in a network of neurons. For example, London et al. (2010) found that adding just a single action potential to a randomly selected neuron in the cortex of a rat produced extra action potentials in 15 to 41 further post-synaptic neurons, thereby leading to a small, but measurable increase in the overall firing rate of the local network. Noise at the level of individual neurons has also been linked to perception and behavior. In one study with monkeys, neural activity in MT, an area of the visual cortex that is responsible for the perception of motion, predicted judgments of motion direction even in “noise trials” when the stimulus had no dominant direction of motion (Britten et al., 1996). It is thus generally agreed that internal noise can rise to consciousness on its own. Or as the psychologist John Wixted (2019, p. 201) puts what he calls the “breakthrough idea” of signal detection theory, “the distribution of sensations generated by neural noise falls above, not below, the threshold of conscious awareness.”

¹ Although I focus on signal detection theory in the main text, noise is presupposed by other prominent models of perception as well, including Bayesian computational models (Kersten et al., 2004, p. 286), sampling models (Hoyer & Hyvärinen 2003), and population codes (Pouget et al., 2000).

Such internal noise is reflected in the phenomenology of experience, as the following description from a perception textbook makes clear.

If you sit in the quietest place you can find and you wear your best noise-canceling headphones, you will find that you can still hear *something*. Similarly, if you close your eyes in a dark room, you still see something—a mottled pattern of gray with occasional brighter flashes. This is internal noise, the static in your nervous system. (Wolfe et al., 2015, p. 12)²

In fact, you needn't be in darkness or silence for the noise in your perceptual system to be reflected in your experience. If you stare at a white wall in good constant lighting you can notice “a subtle ‘flickering’ or ‘crepitation’, of the sort manifest in a snowy, poorly tuned rabbit-ears TV” (Hellie, 2005, p. 493).

2.2 | Noise-Induced Hallucinations

I have just reviewed psychophysical, neural, behavioral, and phenomenological evidence for the existence of perceptual noise. In this subsection, I'll summarize the explanation of mundane hallucinations that I take to be suggested by such noise. New wave relationalists (and disjunctivists) would reject this explanation. But before we consider how they might seek to accommodate mundane hallucinations, it will be helpful to lay out this explanation first.

According to the simple model for thinking about the perceptual experience of low-intensity stimuli, a stimulus is perceptually experienced just in case $I > T$. The existence of perceptual noise renders this model inadequate. Because noise varies randomly from trial to trial, a stimulus of low but constant intensity will be perceptually experienced on some trials but not on others, explaining the S-shaped function depicted in Figure 1b.

Oversimplifying greatly, but harmlessly for current purposes, we can illustrate the influence of noise with a toy example. Suppose that my auditory cortex contains a neuron whose firing rate tracks the presence of a sound. We can think about whether I have a perceptual experience as of a sound as determined by the rate at which this neuron fires. When the neuron's firing rate exceeds the relevant sensory threshold, I have a perceptual experience as of a sound. When it falls below that threshold, I don't. We can conceptualize the neuron's firing rate, in turn, as determined by two factors: the actual intensity of the sound and random fluctuations in the rate at which the neuron fires (i.e., perceptual noise).

We can distinguish two different reasons why the firing rate of my neuron might exceed my sensory threshold: (1) there is a sound that causes a spike in the neuron's firing rate that is sufficient to exceed my threshold; or (2) there is no sound but the neuron's firing rate spikes above my threshold solely as a result of perceptual noise. In (1), I have an auditory experience as of a sound

² The quotation from Wolfe et al. could be interpreted as suggesting that internal noise is the object of your experience—so that you see and hear your own neurons firing. But a less committal interpretation is simply that the internal noise makes a difference to the phenomenal character of your experience, without taking a stand on what (if anything) the object of your experience is. I criticize the idea that we perceive our own neurons in §4.2.

and there really is a sound. My perceptual experience is veridical. In (2), I experience a phantom sound. I have an auditory experience as of a sound, but there is no sound. I hallucinate a sound.³

Of course, whether I perceptually experience a sound is not really determined by the firing rate of a single neuron. This is just a toy example. The important point is that (on the explanation that I'm recommending) whether I have an experience as of a sound is determined by some neural state of mine, which is in turn determined by factors that include both the stimulus intensity and noise. And on some occasions, noise will suffice for the experience as of a sound when the contribution from the stimulus is nil.

Disjunctivists will deny that neural activity of any sort is sufficient for perceptual experience that has the same nature as veridical perceptual experience. While granting that (2) involves a hallucination that is introspectively indistinguishable from (1), they will nevertheless maintain that (1) and (2) differ in phenomenal character (or at least differ in the facts in virtue of which they have their phenomenal characters). For current purposes, I set this response aside.

By contrast, new wave relationalists will deny that (2) involves a perfect hallucination. This commits them to maintaining either that the experience in (2) is introspectively distinguishable from the experience in (1) or else that (2) actually does involve the experience of a mind-independent object. The question, explored in Sections 3–4, is whether either of these claims can be maintained plausibly.

The existence of perceptual noise promises to explain a wide variety of mundane hallucinations. It will be useful to have some examples in mind for the discussion to come (bearing in mind that whether these are *genuine* cases of perfect hallucinations is exactly what's at issue).

PHONE: My cell phone is always set to vibrate when calls are received, and I usually keep it in the front right pocket of my jeans. But every now and then, I have an experience as of my phone vibrating against my upper right thigh even when it isn't in my pocket. A burst of noise—perhaps, a random boost in the firing rates of neurons in my somatosensory cortex—causes me to have the experiences even though nothing is vibrating.

MOSQUITO: While sitting in my backyard on a warm summer night, I was on high alert for mosquitos. At one point, I had an experience as of a mosquito biting my arm. But when I looked, there was nothing there. My experience was a result of perceptual noise.

DOORBELL: While I was in the shower yesterday, I had an experience as of the doorbell ringing. After my shower, I told my wife that I had heard the doorbell while in the shower and asked her who was at the door. Confused, she replied that the doorbell didn't ring. My experience had been caused by perceptual noise.

CRY: When I was a new parent, I would sometimes have experiences as of my child crying out from his bed in the silence of night while I was down the hall in my own bedroom. But often no one had cried out. The child was fast asleep. My experiences had been caused by the random firing of neurons in my auditory system.

³ For the sake of simplicity, I set aside two other cases that aren't immediately relevant: (3) there is a very faint sound and a modest burst of noise such that the sound and noise are jointly sufficient, but individually insufficient, to cause a spike in the neuron's firing rate that exceeds my threshold; (4) there is a sound and a large burst of noise, either of which is individually sufficient to cause the neuron's firing rate to fire above threshold. Case (4) is an instance of overdetermination, and presumably counts as veridical. Case (3) is harder to classify. On the one hand, my perceptual experience as of a sound is lucky; had it not been for a burst in random noise I wouldn't have had it. So, one might be tempted to classify it as illusory or hallucinatory. But on the other hand, I have a perceptual experience as of a sound and there really is a sound. Moreover, the sound is a cause of my perceptual experience. One might thus be tempted to classify my perceptual experience as veridical. This conclusion is encouraged by the hypothesis that a key function of noise is to enable the detection of faint stimuli (Wiesenfeld & Moss 1995; McDonnell & Ward 2011).

STARS: On a warm, clear summer night, I was lying in a field looking up at the sky with my kids. Having just taught them about shooting stars, we were all staring upwards in search of them. At one point, I declared that I had just seen one in my periphery. But no one else had seen it because it didn't exist. My experience had been caused by photoreceptor noise; rhodopsin molecules in my eye activated spontaneously.

Note that although I describe each of these examples as being generated by internal perceptual noise, I do not mean to imply that controlled studies have confirmed this. For example, I don't know of any controlled studies subjecting people to hungry mosquitos while scanning their brains to see if, when, and how false alarms arise. But when researchers have used stimuli that better lend themselves to a laboratory setting, they have found that perceptual noise (whether internal or external) can bring about hallucinations—for example, of flashes of light (Barlow, 1956), oriented gratings (Pajani et al., 2015), and the letter S (Gosselin & Schyns, 2003). Taken in conjunction with the broad empirical reasons, canvassed above, to think that internal perceptual noise is a real phenomenon that modulates conscious experience, these findings are strongly suggestive of the hypothesis that mundane hallucinations are likewise caused by perceptual noise, including sometimes internal noise. I am not the first to endorse this hypothesis. For example, Wixted (2019, pp. 202, 212) proposes that phantom phone vibrations are caused by neural noise.

Going forward, my working hypothesis will thus be that mundane hallucinations caused by external and internal noise occur in tactual, auditory, and visual perception, and are commonplace. (How commonplace? One indication comes from Drouin et al. (2012), who report that 89 percent of undergraduates experience phantom phone vibrations, on average once every other week.) On the face of it, this creates problems for the new wave relationalist view that object-less experiences that are indiscriminable from veridical perceptual experiences are impossible, for here we have what appear to be actual examples of such experiences.

Although I've been concentrating on mundane hallucinations, other hallucinations have also been explained in terms of perceptual noise. For example, visual hallucinations associated with Charles Bonnet Syndrome (ffytche, 2005) and auditory hallucinations associated with tinnitus (Langguth et al., 2013) often emerge following the degeneration of input pathways, and it is thought that central neurons compensate by amplifying their firing in response to all inputs, making them more sensitive to noise. Similarly, one explanation of the auditory verbal hallucinations experienced by schizophrenics appeals to spontaneous activity in auditory areas of the brain (Cho & Wu, 2013). While one might try to leverage these pathological hallucinations against new wave relationalism, doing so would present unique challenges, and considerably lengthen the current treatment. I'll thus remain focused on mundane hallucinations here.

In the next two sections, I'll argue that new wave relationalists' usual strategies are not available to explain away mundane hallucinations.

3 | THE DENIAL STRATEGY

The new wave relationalist's first line of defense is to deny that putative perfect hallucinations are indiscriminable from veridical perceptual experiences. This section evaluates three ways of developing this denial strategy.

3.1 | Denying that mundane hallucinations are possible

The first approach we can immediately set aside. In response to many of the more outlandish extraordinary hallucinations, new wave relationalists often express skepticism that they are genuinely possible. For example, in a scenario that he calls the “chaos hypothesis,” Chalmers (2005, p. 23) imagines that “random, uncaused experiences” emerge due to some quantum or other freak accident that “are exactly the sort of regular, structured experiences with which I am familiar.” Likewise, Papineau (2014, p. 8) imagines “a ‘cosmic swampbrain’ that has randomly assembled itself along with supporting vat in outer space.” In his defense of new wave relationalism, Raleigh (2014, p. 104) replies,

One might try to deny that such a chaotic case is really possible — or, less ambitiously, one might try to insist that such a ‘merely logical’ possibility is not something that a theory of consciousness needs to deal with.

Similarly, in reply to Kriegel’s (2013) imagined scenario of a “space soul,” a disembodied mind floating through space that experiences hallucinations that are phenomenally just like our perceptual experiences, Masrour (2020, p. 757) objects that absent positive evidence for their actuality or empirical possibility, “the conceivability of space souls does not provide any independent evidence for their possibility.” Even more strongly, Noë (2007, p. 468) writes, “It strikes me as crazy to think that we now know that a disembodied, dangling, accidental brain would have experiences qualitatively indistinguishable from mine” and “We have no reason to think there is any experience going on here at all.”

There is, however, one very important respect in which mundane hallucinations are unlike chaotic hallucinations, swampbrain hallucinations, or space soul hallucinations. Mundane hallucinations are actual. Their possibility is thus not in question. If new wave relationalists have a promising way to address them, it will thus need to derive from another approach—one which challenges not the possibility of such experiences, but rather their correct characterization.

3.2 | Denying that mundane hallucinations have perceptual phenomenology

New wave relationalists are not committed to the view that all experiences have mind-independent objects. They are only committed to the view that all *perceptual* experiences have mind-independent objects. In fact, relationalists are often motivated by the idea that perceptual experiences have a special phenomenal character that other experiences lack. It is thus open to new wave relationalists to classify certain alleged hallucinations as non-perceptual experiences. For example, Masrour (2020, pp. 750–752) denies that dreams are perceptual experiences. Instead, he classifies them as pseudo-perceptual experiences, which may be similar to perceptual experiences in some respects, but lack essential aspects of the phenomenal character that defines perceptual experiences. Whereas perceptual experiences are, perhaps, marked by qualities such as constancy, presence, vividness, saturation, and/or determinacy, pseudo-perceptual experiences lack one or more of these qualities, and are thus in-principle discriminable from perceptual experiences (Masrour, 2020, p. 751). Masrour argues that we lack good evidence that dreams are perceptual experiences rather than pseudo-perceptual experiences. Similarly, Noë (2004, p. 214; 2007,

p. 471) claims that dreams lack the phenomenal detail, phenomenal stability, and sense of involvement of perceptual experiences. And Raleigh (2014, p. 85) claims that we have the phrase “dream-like” precisely because dreams differ phenomenally from perceptual experiences.

While some people might report that their dreams are just like their perceptual experiences, there are reasons to question the reliability of such reports. For instance, dream reports require dreamers to correctly remember their experiences, but memory may play a distorting role. There are also significant differences in one’s overall neural activity during dreaming and wakefulness. Masrour is thus skeptical that dreams are perceptual experiences. Moreover, Masrour proposes to extend this strategy beyond dreams to “all cases of internally generated experiences that are allegedly phenomenally indistinguishable from veridical perceptual experiences, including experiences resulting from psychosis or hallucinogens, iconic memories, and some cases of imagination” (2020, p. 750). So perhaps he would also attempt to extend the strategy to mundane hallucinations that are generated by internal noise.

There are, however, two obstacles to extending the strategy in this way. The first is that reports of mundane hallucinations are not subject to the same worries about reliability as reports of many other experiences. You can never report a dream while it is happening. Dream reports are thus always filtered through memory. But you can report a mundane hallucination in real time. Reports of experiences caused by psychosis and hallucinogens don’t rely on memory, but their reliability is also questionable insofar as the reporter’s domain-general mental capacities are compromised. Mundane hallucinations, by contrast, occur to everyone, including those who are clean, sober, and mentally fit. It is thus much harder to plausibly discount the reports of those who claim to have mundane hallucinations that are phenomenally identical to their veridical perceptual experiences. When I tell my wife that I just heard our son yell “Mama!” from his room, she can fairly tell me that I am wrong, that our son is in fact fast asleep. But she cannot fairly accuse me of misremembering my own experience, let alone of being high or psychotic. Of course, I could be mistaken all the same. Reports of experiences can be influenced by decision or response biases. But as Masrour acknowledges, the bar is not supposed to be certainty. It’s supposed to be “good evidence according to the evaluative criteria that are implicit in the practice of empirical science” (2020, p. 752). And all things equal, the empirical science takes the honest reports of sober, attentive, mentally healthy individuals to be good, if defeasible, evidence about the phenomenal character of their current experiences.

It is, moreover, often possible to empirically corroborate first-person reports and to test for influences from bias. For example, Powers et al. (2017) induced hallucinations through Pavlovian conditioning. After a visual stimulus was repeatedly paired with a tone, subjects reported hearing the tone when only the visual stimulus was presented. But wait! Did the subjects really *hallucinate* the tone? Or did they merely *judge* that the tone must have been present given that the visual stimulus was present (a form of decision bias)? Using fMRI, the researchers isolated areas of the auditory cortex that respond to actual tone presentations, and that generate clear auditory experiences when electrically stimulated. They then found that activity in those areas during post-conditioning trials in which only the visual stimulus was presented correlated with reports of hearing phantom tones. When the activity in those areas was high, subjects reported hearing a tone. When the activity was low, subjects reported hearing no tone. The researchers conclude that the “conditioned hallucinations involved true percepts” (p. 597).

There is a second problem with extending new wave relationalism’s treatment of dreams to mundane hallucinations. One’s overall neural activity is known to be very different during dreams and waking life. This adds plausibility to the hypothesis that dreams differ phenomenally from waking experiences. But when mundane hallucinations are produced by internal noise, there are

no comparable known neural differences between them and veridical experiences. To return to the toy example from Section 2, the neuron in my auditory cortex that tracks the presence of sound can increase its firing rate because there is a sound or because of random noise. The difference in the two cases isn't a neural difference, but an etiological one. You have to look outside of the brain to find out whether the boost in neural firing was caused by a stimulus. Of course, new wave relationalists are phenomenal externalists. They allow that two neurally identical experiences can differ in phenomenal character when one experience has a mind-independent object and the other does not. But the point is not that new wave relationalists are *forced* to admit that internal noise-induced hallucinations are phenomenally identical to veridical perceptions. Rather, the point is that when neural differences are known to exist between veridical experiences and alleged hallucinations (as they do, for example, between veridical waking experiences and dreams), those differences bolster the contention that the two experiences differ phenomenally. But when there are no known neural differences, the contention cannot be bolstered in the same way. Thus, while new wave relationalists can claim that mundane hallucinations are discriminable from veridical experiences, it is unclear why anyone not independently committed to relationalism should believe them.

New wave relationalists might object that the toy example is hiding potential neural differences between veridical and noise-induced experiences. Veridical perception involves a cascade of neural events. Even if there is noise that increases the firing rate of one or more neurons in my auditory cortex, there may be other (e.g., sub-cortical) neurons whose firing patterns differ between veridical and noise-induced experiences. Perhaps those other neurons are essential to making an experience perceptual rather than pseudo-perceptual, even if the phenomenal difference is subtle enough that we do not ordinarily notice it introspectively. For example, maybe those other neurons are necessary to give rise to the quality of constancy (or vividness, or...) that Masrour takes to be possibly constitutive of perceptual experience.

But this objection sits uncomfortably with the empirical evidence, cited earlier, that a single extra action potential can itself lead to a cascade of additional neurons firing, and that activity in MT predicts perceptual judgments of direction of motion even on noise trials. It also overlooks the fact that internal noise can enter at any stage of perceptual processing, including at the earliest receptor stages where fluctuating membrane potentials or the random activation of photoreceptors could trigger a cascade through the same neural mechanisms (including constancy mechanisms) that are involved in the processing of a real signal. When the noise enters at such an early stage, as in STARS, the full neural profile is likely to be no different than if it had been caused by an external stimulus of the appropriate type. To appreciate this point, it is important to recognize that the visual system is extremely sensitive. Human seers can detect a *single photon* above chance (Tinsley et al., 2016). This detection occurs because rhodopsin (a photopigment in the eye) isomerizes in response to the photon. But as we have seen, rhodopsin also isomerizes spontaneously. So a given photoreceptor might fire either in response to a photon or because of random noise. In either case, the downstream neural consequences will be the same.

While I have focused on a single photon and a single photoreceptor to make the point vivid, in principle noise can also cause multiple photoreceptors (or other cells) to fire, and thus mimic the neural profile of a more complex stimulus, though the probability of noise exactly matching the pattern of a stimulus will decrease as the stimulus's complexity increases. In practice, many stimuli will be sufficiently complex that the probability of a neural match arising from noise will be so low that it will never occur. There is a reason that mundane hallucinations almost always mimic experiences of very simple stimuli.

In conclusion, it is hard to see how new wave relationalists can extend their treatment of dreams to mundane hallucinations. While there are known neural differences between dreams and waking perceptual experiences that lend independent support to the introspectively plausible suggestion that dreams and waking experiences differ phenomenally, there is (so far as I am aware) no known neural difference between mundane hallucinations and veridical perceptual experiences, some reason to think that no neural difference exists (especially when the noise arises at the earliest receptor stages), and no compelling introspective grounds to think that mundane hallucinations and veridical perceptual experiences differ phenomenally.

3.3 | Denying that mundane hallucinations have sufficient temporal extension

There is one further way of maintaining that mundane hallucinations are discriminable from veridical perceptual experiences that deserves consideration. According to Raleigh (2014, p. 96), “A relational theorist should be committed to the idea that a perceptual experience is not instantaneous—it always has some non-negligible duration.” Otherwise, “the subject does not achieve perceptual contact with the external world, and instead enjoys mere visual sensation.” Thus, it might be claimed that genuine perceptual experiences cannot be ephemeral; they need to span a certain duration. Since mundane hallucinations are generally fleeting, they would be discriminable from perceptual experiences on that basis.

Raleigh does not specify how long an experience needs to last to count as having a “non-negligible duration,” and he never discusses mundane hallucinations, so it’s not clear that he would take this temporal extension requirement to disqualify them. But one could imagine a new wave relationalist adapting Raleigh’s requirement towards this end.

Notice, however, that any form of the temporal extension requirement that would disqualify mundane hallucinations from being perceptual experiences would also disqualify many veridical experiences from counting as perceptual experiences. It would mean that we never see shooting stars, a flash of lightning, or a briefly met gaze; hear a sharp knock at the door, a clap of thunder, or a child cry out “Mama!”; or feel a mosquito bite, a sudden breeze, or a peck on the cheek. Many of our veridical experiences are short in duration, and the relationalist would have to insist that all of them are “mere sensations” that involve no “perceptual contact with the external world.” Not only is that implausible on its face, but it undercuts a central motivation for relationalism, which is to put perceivers in touch with the world in a way that representational and sense-data theories of perception allegedly do not.

It is also unclear what is supposed to motivate the temporal extension requirement. Why should the *duration* of an event matter as to whether it is perceived? It’s true that an extended event is necessary to support exploratory activities, such as reaching and manipulation, which are key functions of perception. But it certainly doesn’t follow that perception *requires* “the right kind of ‘exploratory’ activity/impulses on the subject’s part” (Raleigh, 2014, p. 96). After all, another key function of perception is detecting stimuli at a distance, when they’re beyond the range of exploration. Some stimuli are best left unexplored. The appeal to the temporal extension requirement thus appears ad hoc.

But suppose we grant the temporal extension requirement for the sake of argument, and maintain that an experience needs to surpass duration D in order to be perceptual. Suppose too that the longest known duration of any actual mundane hallucinations is less than that, $D - \epsilon$. The new wave relationalist would then need to deny that perceptual noise could generate an experience of

duration D . But once one grants that perceptual noise *actually* generates experiences of duration $D - \epsilon$, it's hard to see why we would deny that it is metaphysically possible for perceptual noise to generate experiences that are only ϵ longer in duration. The known existence of mundane hallucinations thus opens the door to powerful conceivability arguments for the possibility of more sustained hallucinations—conceivability arguments that do not require the sort of heavy-duty background assumptions that are needed to support the conceivability of swampbrains or space souls.

4 | THE REDESCRIPTION STRATEGY

Section 3 criticized attempts to deny that mundane hallucinations are indiscriminable from veridical perceptual experiences. I argued that new wave relationalists cannot plausibly claim that mundane hallucinations are impossible, have only a pseudo-perceptual phenomenology, or are too short to qualify as genuine perceptual experiences. But this still leaves the redescription strategy unexplored. In this section I consider whether mundane hallucinations can be redescribed as illusions, experiences of non-standard objects, or experiences of past objects.

4.1 | Illusionism

Perceptual experience involves the attribution of properties to objects. A perceptual illusion occurs when a property is attributed to an object though the object does not have that property. The failure is thus one of attribution. The stick in water looks to be bent, but it is actually straight. A hallucination, by contrast, occurs when there is a failure of reference. A property is attributed to an object, but there is no object.

One version of the redescription strategy redescribes hallucinations as illusions. The drug user doesn't hallucinate an elephant in his room; he illusorily perceives the room as containing an elephant. In this way, a failure to be properly related to an object is transformed into a failure to properly attribute a property. This strategy is championed by Ali (2018), who calls it "illusionism," and Masrour (2020), who calls it "displacement." Can it be applied to mundane hallucinations?

One *prima facie* difficulty is that it isn't always clear what could plausibly serve as a perceived object in the case of mundane hallucinations. The drug user may be wrong about the existence of the elephant, but he veridically perceives the walls of the room, and so the room is available to serve as the object of his experience. Some mundane hallucinations are like this. In DOORBELL, I "hear" the doorbell against the backdrop of the sound of the shower, which I perceive veridically. It might thus be argued that my experience is an illusion. But for other mundane hallucinations, it is less obvious what could serve as the object of experience. In CRY, I "hear" a cry when it would seem that I'm not hearing anything else. Thus, one might worry that for mundane hallucinations there isn't always an object available to support a redescription in terms of an illusion.

But perhaps we just need to be more creative. According to Langsam (1997), some visual hallucinations are relations to empty space. Similarly, Alston (1999) briefly entertains the hypothesis that Macbeth is related to the air in front of his face when he hallucinates a dagger. And Ali (2018) conjectures that when you hallucinate with your eyes shut, you're related to your eye lids; and when you hallucinate in a void, you're related to the void. Drawing inspiration from such proposals, we might search for an object to serve in mundane hallucinations. For example, in CRY we might say that I perceive a region of space down the hall as containing a child's cry.

Of course, this suggestion is controversial. Not everyone thinks that we ordinarily hear regions of space. Here is one reason to be suspicious of redescribing hallucinations as illusions in this way. A plausible necessary condition on an entity's serving as the object of a perceiver's perceptual experience is that the perceiver be causally related to the object. This is true not only in veridical perception, but also in illusion. You are causally related to the stick even if you perceptually experience it as bent. In some of the cases we've considered, this necessary condition is satisfied. The drug user is causally related to the room when he has an experience as of an elephant in a room. But it is doubtful that I am causally related to the relevant regions of space during the experience in CRY. In that case, the experience is produced endogenously, solely through internal noise.

Another problem with redescribing mundane hallucinations as illusions is that it often conflicts with the phenomenology. When you hear your child's cry, the cry presents itself as the object of your experience. It is the cry that draws your attention, not the region of space from which it derives (you might not even be sure where it is coming from). And the same is true when you hallucinate a cry. Admittedly, claims about phenomenology rarely command universal agreement, and if a new wave relationalist claimed that CRY and other mundane hallucinations are best described in terms of regions of space being the object of experience, I'm not sure I would have anything much to say to persuade her otherwise. But one reason to be skeptical that this is the natural, pre-theoretical view, is that it conflicts with standard accounts of the experience of sounds—accounts that were developed independently of the present debate. Thus, according to O'Callaghan (2007), sounds are individuals, events in which a medium is disturbed, and to have an auditory experience as of a sound is, in the usual case, to have an experience as of such an individual. While the new wave relationalist could reject this account, she would then incur a debt to defend an alternative.

4.2 | Anchoring

In his commentary on the film *The Matrix*, Chalmers (2005) argues that a brain in a vat whose experiences were controlled by a computer simulation would not hallucinate an external world. Rather, it would veridically perceive a virtual world. In Masrour's (2020, pp. 752–756) terminology, its experiences would be *anchored* to the computer that is causally controlling its experiences; some part of the computer would serve as the object of its experiences. Raleigh (2014) and Masrour (2020) adopt variations on this anchoring strategy to redescribe the experiences of envatted brains, evil demon victims, and others whose experiences have a deviant external cause.

On the face of it, however, this strategy seems ill suited to redescribe mundane hallucinations that are caused by internal perceptual noise. Such hallucinations are the product of random fluctuations in biological activity, such as neural firing rates, and are thus generated endogenously. They are not controlled by a computer, evil demon, or anything else external to the mind.

In reply, it might be suggested that the strategy of anchoring shouldn't be restricted to external causes. Thus, new wave relationalists might seek to anchor mundane hallucinations in the perceiver's internal states. This proposal needs to be handled with care. On pain of collapsing into a sense-data theory, such internal states shouldn't be identified with internal *mental* objects. But they might be identified with internal *neural* objects, so long as those specific neural objects are not themselves identical to mental objects. So, if my experience in PHONE is caused by the spontaneous activity of some neurons in my somatosensory cortex, it might be claimed that those neurons are the object of my experience. Of course, the new wave relationalist can allow that those neurons are perceived illusorily. They seem to have the property of vibrating at the same rate that

my cellphone usually vibrates at, when in fact they aren't vibrating at all. But still, my experience attributes that property to the neurons in my somatosensory cortex.

While this is an available position in logical space, extending the anchoring strategy to internal neural objects in this way doesn't appear to be terribly popular. Alston (1999, p. 191) mentions the possibility that we perceive our own brain states in one sentence, but deems it implausible and immediately sets it aside. The only other relevant discussion I'm aware of is due to Ali (2018, p. 622), who says that "phosphenes and related phenomena, are internal objects, or at least events that internal objects undergo" and that they "are features of our perceptual systems, and this makes them readily available as objects of perception." But Ali doesn't specify what kinds of features or events he has in mind, and doesn't explicitly discuss noise-induced hallucinations (phosphenes are most commonly caused by pressure on the eyeballs), so it's not clear that he would endorse this position either.

In evaluating the anchoring strategy, it is important to distinguish the (plausible) thesis that the object of a perceptual experience is always a cause of that experience from the (implausible) thesis that any cause of an experience is its object. If an apple drops from a tree onto the back of your head, coincidentally causing you to have a visual experience as of an apple, you do not thereby *see* the apple. That X caused your experience does not entail that X is an object of your experience. So, if the new wave relationalist wants to endorse the anchoring strategy, she must do more than point out that your neurons cause your mundane hallucinations. She must further explain why your neurons are the objects of your mundane hallucinations. To do this, she might reach for an off-the-shelf naturalized theory of perceptual content. But a roadblock awaits. The best such theories have a teleosemantic component; they appeal to *functions* to help determine content (Dretske, 1988; Millikan, 1984; Neander, 2017; Shea, 2018). But given that a response to pure noise is generally a *malfunction*, it is hard to envision a convincing story according to which perceptual experiences have the function to refer to neural states engendered by pure noise. The path forward for new wave relationalists who want to argue that neural states are the objects of mundane hallucinations is thus far from straightforward.

A further problem with the anchoring strategy is that it is in tension with the favored idea of relationalists that experience is "transparent" to the external world. This is typically taken to mean not merely that experience provides direct phenomenal and epistemic access to the world that is external to our *minds*, but also that it provides direct access to the world that is external to our *heads*—the world of mulberry trees, lampposts, and nosy neighbors. But if the phenomenal character of veridically hearing your child's cry is the same as the phenomenal character of hallucinating your child's cry, and if hallucinating your child's cry is a matter of perceiving your own neurons, then it is hard to see how new wave relationalists can claim that the world beyond your head is phenomenally present in perceptual experience. Similarly, if hearing and hallucinating your child's cry are phenomenally identical, with the latter involving a relation to your neurons, it's hard to see how veridical experience could provide direct epistemic access to the world beyond your head in virtue of its phenomenal character.

A final problem with extending the anchoring strategy to internal states takes the form of a dilemma. On the one hand, new wave relationalists can maintain that one's neurons are only the object of experience when there is no external object. In that case, however, they are obliged to explain why the presence or absence of an external object should make a difference to whether one perceives one's own neurons. If I perceive my own neurons when I hallucinate my phone vibrating, why don't I also perceive them when I veridically feel my phone vibrating? One possible answer, suggested by Ali's (2018, pp. 622–624) treatment of phosphenes, is that my neurons aren't salient when I'm attending to external objects. They only become salient when I have nothing

external to attend to. But attention can be endogenously controlled, and even non-salient stimuli can be attended to with effort. For example, while a flashing light is more salient than a white wall, I can force myself to attend to the wall rather than the light. So, when my phone vibrates, I likewise ought to be able to shift my attention back and forth between the phone and my neurons. But (speaking for myself) this is not something I seem to be able to do.

On the other hand, new wave relationalists can maintain that we constantly perceive our own internal states. Right now, you not only perceive the words on this page; you also perceive your own neurons. This view is certainly counterintuitive. Most people do not take themselves to be constantly perceiving their own neurons. And while some philosophers are happy to flout commonsense, relationalists often defend their view under the banner of preserving it. Furthermore, this view suggests that perception is typically illusory, since if we do constantly perceive our own neurons, we would seem to perceive them as having properties that they do not actually have. For example, when I veridically perceive my phone vibrating, I have an experience as of a vibration of a certain intensity; but my neurons are not vibrating at that intensity.

I conclude that anchoring, like illusionism, is not a promising approach to explain away mundane hallucinations.

4.3 | Past Objects

When I recall my eldest son's first steps, my son is an object of my memory. Perceptual experiences are thus not the only experiences that have objects. Episodic memories also have objects (Barkasi & Rosen, 2020). If hallucinations had objects supplied by memories, then perhaps that could explain why they are indiscriminable from veridical experiences. In this spirit, philosophers have proposed that some (Barkasi, 2020) or all (Byrne & Manzotti draft) hallucinations involve the awareness of an object that is retrieved from memory.

For certain hallucinations, this proposal enjoys some intuitive plausibility. If, like Hamlet, you hallucinate your father, it would not be unreasonable to suppose that your memory of your father is somehow involved. Or consider Albright's (2012, p. 238) example of a magician who induces the experience of seeming to see a ball vanish (as discussed by Barkasi, 2020, pp. 466–467). After repeatedly throwing a ball in the air and catching it, the magician feigns a final toss, holding onto the ball but following its would-be trajectory with his gaze. To the audience, it looks as though the ball disappears mid-flight. A not wholly outrageous hypothesis is that the object of the hallucination is supplied by a memory of the very ball that the audience has just seen repeatedly tossed.

When we consider mundane hallucinations, however, this hypothesis seems far less plausible. For one thing, there is no single object that emerges as the most likely candidate of any given mundane hallucination. In PHONE, for example, which of the thousands of phone vibrations that I've felt in the past is supposed to be the object of my experience? Note that the problem is not merely that it's hard to know which past vibration I'm recalling. It's worse than that. Because my current hallucination does not plausibly have any single previously experienced vibration as its dominant causal source, there doesn't seem to be any fact that would make it the case that I'm recalling one vibration rather than another.

A further problem is that recalled objects are not usually indiscriminable from presently perceived objects. When I recall seeing my eldest son take his first steps, I am not under the impression that I am *currently* seeing him walk. Likewise, when I remember how he would yell "Dada!" from his crib, I do not infer that he is *right now* calling for me. We are, in short, usually pretty good

at distinguishing experiences of past and present objects. So anyone who maintains that hallucinations essentially involve the awareness of past objects is indebted to explain why they aren't recognized as such. For some hallucinations, such an explanation is not too hard to conjure. When Hamlet hallucinates his father, he's not at his best. Maybe his melancholy has something to do with his conflation of past and present. But if mundane hallucinations involve recalled objects, what is the explanation of why they aren't recognized as such? It is, at the least, non-obvious.

Finally, insofar as hallucinating involves remembering, patterns of neural activity that are associated with memory ought to be implicated in hallucination. For some hallucinations, this may be the case. For example, Albright (2012) speculates that the vanishing ball illusion involves top-down processing that is consistent with the deployment of mental imagery. But it isn't true of mundane hallucinations that are induced by perceptual noise. At a neural level, such noise is associated with perceptual pathways. In STARS, for example, the hallucination is generated in a bottom-up manner, when rhodopsin molecules in the retina isomerize spontaneously. The hypothesis that mundane hallucinations have recalled objects thus seems to be at odds with what we know about their neural realization.

Redescribing mundane hallucinations as relations to past objects is thus no more likely to save new wave relationalism than is redescribing them as illusions or relations to non-obvious objects. The redescription strategy lacks promise.

5 | THE ARGUMENT FROM HALLUCINATION REVISITED

New wave relationalism is motivated by the argument from hallucination, which requires an appeal to *perfect* hallucinations—that is, experiences which lack an object but are introspectively indistinguishable from veridical experiences. Many of the experiences that we ordinarily call “hallucinations” are not like this. When you take LSD and a houseplant looks to you like a dragon, we might say that you “hallucinate” a dragon. But since you really illusorily perceive a houseplant, your experience has an object. Thus, your experience cannot figure into the argument from hallucination to put pressure on relationalism. Likewise, if someone with Charles Bonnet Syndrome hallucinates cartoonish figures that are obviously unreal, those hallucinations cannot support a plausible version of the argument from hallucination. Given their introspective discriminability from veridical experiences, it is easy to deny that they have the same nature as veridical perceptual experiences.

The argument from hallucination thus requires perfect hallucinations. But the perfect hallucinations that philosophers typically summon—brains in vats, swampbrains, space souls, and the like—are products of science fiction. They require a leap of imagination that leaves room for doubt about whether they are truly possible. One might therefore be inclined to follow Allen's advice that they shouldn't drive theorizing.

Any philosophical theory of hallucination will need to say something about purely hypothetical ‘philosophical’ hallucinations... However, the philosophical obsession with merely possible cases of hallucination might seem methodologically questionable. In particular, why should so much weight be given to intuitions about remote counterfactual situations in deciding between theories of hallucination, and correlatively between theories of perception? Instead of treating purely hypothetical cases as central to theorising about the nature of hallucination, and thereby perception, an alternative approach is to focus in the first instance on actual instances of the

phenomenon, and to treat purely hypothetical cases of hallucination as ‘spoils for the victor’: as theoretically marginal cases to be accommodated by an independently motivated theory. (Allen, 2015, pp. 290–291)

Although Allen is defending a version of disjunctivism, his remarks here equally capture the perspective of many new wave relationalists. But crucially, this line of thought cannot be extended to mundane hallucinations since they are not hypothetical. They are nature’s perfect hallucinations.

Mundane hallucinations thus strengthen the argument from hallucination because they are actual. They cannot be dismissed as philosophers’ fantasies. They may strengthen the argument from hallucination for another reason too. Martin (2004) has argued that the strongest versions of the argument from hallucination involve *causally matching hallucinations*, which have the same proximal, or neural, cause as veridical perceptual experiences (see also Robinson, 1994). There are two reasons for this.

First, causally matching hallucinations are especially strong candidates to serve as perfect hallucinations. As a general empirical matter, there are robust correlations between neural activity and reported experiences, and while those correlations are perhaps insufficient to prove that causally matching hallucinations are indiscriminable from veridical experiences, they could be used to support an abductive argument for that conclusion. All things equal, the relationalist should thus prefer to avoid commitment to the thesis that causally matching hallucinations are discriminable from veridical perceptual experiences. Thus, if mundane hallucinations are causally matching hallucinations, that would lend further support to the premise that perfect hallucinations are possible.

Second, causally matching hallucinations can lay an especially strong claim to having a nature that is shared by veridical perceptual experiences, lending support to the other key premise in the argument from hallucination. That is because the nature of a causally matching hallucination would seemingly need to consist in properties that are internal to the thinker, and those internal properties are shared when the perceiver has a corresponding veridical perceptual experience. This in turn leads to a screening off worry: if the internal properties suffice to explain the phenomenal character of the causally matching hallucination, it’s hard to see why they wouldn’t also suffice to explain the phenomenal character of the corresponding veridical experience. Of course, there are things that disjunctivists can say here to try to address the screening off worry, and Martin (2004) himself recommends a line that he thinks disjunctivists should take. But the point is that the screening off worry forces disjunctivists to say *something*. Causally matching hallucinations thus make the argument from hallucination stronger than it would otherwise be.

So, if mundane hallucinations were causally matching hallucinations, that would make them especially powerful. Are they? To address this question, we first need to get a little clearer on what is required for a mundane hallucination to *causally match* a veridical experience. One thing that would *not* be required is strict neural identity throughout the entire brain. Instead, the match should be limited to those aspects of neural activity that are responsible for the experience. Here, we can draw on the concept of a *neural correlate of consciousness*, understood as the minimal subset of neural states that are jointly sufficient for a given conscious experience (Chalmers, 2000). The question whether mundane hallucinations are causally matching hallucinations thus morphs into the question whether the neural correlates of veridical experiences and mundane hallucinations are identical.

One reason to think that they sometimes are is that mundane hallucinations can be caused by perceptual noise that matches the upstream causes of veridical perceptual experiences—the

causes that immediately precede their neural correlates. To see this, suppose that rhodopsin isomerizes, activating a photoreceptor in the retina. By all accounts, the neural correlates of visual experience are not located in the retina, but rather downstream of the eye. This downstream activity should be the same whether the isomerization of rhodopsin was caused by a photon or mere noise. After all, the downstream neural events have no way of knowing *why* the photoreceptors activated. To suppose otherwise would be to commit to a kind of “action at a distance in the neurological realm,” a hypothesis that “few neuroscientists with the hope of serious funding would pursue” (Martin, 2004, p. 54). The neural correlate of the experience in STARS should thus match that of its veridical counterpart.

Some of the other examples of mundane hallucinations may be less clear cut. For example, suppose that my experience in CRY is caused by the random firing of neurons in my auditory cortex—a possibility that is consistent with the fact that direct electrical stimulation of neurons in auditory cortex elicits auditory experiences (Penfield & Perot 1963). Does that pattern of firing exactly match the pattern of firing that constitutes or immediately causes the neural correlate of my veridical experience of hearing my child call out in the middle of the night? The answer would depend on the nature and locus of the neural noise, and we don’t have the same detailed understanding of that noise as we do of photoreceptor noise. But it at least seems nomologically possible for the experience to be causally matching. All it would take is for neural noise to arise in the right parts of the brain. And since neural noise arises from stochastic sub-cellular events that are inherent to all neurons, such as channel noise and the spontaneous release of neurotransmitter vesicles, it seems to be nomologically possible for neural noise to arise anywhere in the brain. Without appealing to exotic science-fiction scenarios, we thus have a compelling story to tell about how CRY, or any of the other mundane hallucinations we have considered, could be causally matching hallucinations. Mundane hallucinations are thus once again very different from the fantastical hallucinations that relationalists typically consider.

6 | CONCLUSION

In focusing on extraordinary hallucinations, new wave relationalists have overlooked mundane hallucinations. And I have argued that their usual strategies of denial and redescription cannot be plausibly extended to these. While there are surely places where new wave relationalists might push back or dig in, I have tried to indicate how significant costs are likely to be incurred as a result. The upshot is that it is very hard to see how new wave relationalism can plausibly respond to the argument from hallucination without falling into disjunctivism. If that is right, then new wave relationalists may wish to reconsider their options.

On the one hand, they might maintain their commitment to relationalism by reembracing disjunctivism. Even here, however, mundane hallucinations introduce new complications. For while it is one thing to maintain that *extraordinary* hallucinations and veridical perceptual experiences have different natures, it is something else to maintain that *mundane* hallucinations and veridical perceptual experiences have different natures. Signal detection theory treats noise as an ineliminable and non-pathological component of all perceptual processing, and gives the same basic explanation of how hallucinatory and veridical experiences arise (Wixted 2009). Moreover, noise is often put to good use by the perceptual system. For example, when visual and haptic cues provide conflicting information about the height of a stimulus, the perceptual system resolves the conflict, in part, by privileging the more reliable cue; and when noise is introduced into one of the cues, that cue is given less weight (Ernst & Banks, 2002). Noise can be beneficial in other

ways too. It can help perceivers detect a signal that would otherwise be too faint—a phenomenon known as *stochastic resonance* (Wiesenfeld & Moss, 1995; McDonnell & Ward, 2011). It may also increase the *dynamic range* of neurons—that is, the range of stimulus intensity values to which they are sensitive (Shadlen & Newsome, 1998). But if noise is ubiquitous and beneficial in these ways, then it might be argued that the assumption that mundane hallucinations and veridical perceptual experiences share a common nature coheres better with the science of perception than its denial—though whether such an argument could be successfully developed is not something I will undertake to determine here.

On the other hand, these considerations, as well as other concerns about disjunctivism, might persuade new wave relationalists to abandon relationalism altogether. Which option is best, disjunctivism or anti-relationalism, is the subject of a very familiar debate; and, if the arguments presented here are correct, it is a debate that is not easily avoided.⁴

REFERENCES

- Albright, T. D. (2012). On the perception of probable things: Neural substrates of associative memory, imagery, and perception. *Neuron*, 74, 227–245. <https://doi.org/10.1016/j.neuron.2012.04.001>
- Ali, R. (2018). Does hallucinating involve perceiving? *Philosophical Studies*, 175, 60–627. <https://doi.org/10.1007/s11098-017-0884-7>
- Allen, K. (2015). Hallucination and imagination. *Australasian Journal of Philosophy*, 93, 287–302. <https://doi.org/10.1080/00048402.2014.984312>
- Alston, W. P. (1999). Back to the theory of appearing. *Noûs*, 33, 181–203. <https://doi.org/10.1111/0029-4624.33.s13.9>
- Ashmore, J. F., & Falk, G. (1977). Dark noise in retinal bipolar cells and stability of rhodopsin in rods. *Nature*, 270, 69–71. <https://doi.org/10.1038/270069a0>
- Barkasi, M. (2020). Some hallucinations are experiences of the past. *Pacific Philosophical Quarterly*, 101, 454–488. <https://doi.org/10.1111/papq.12320>
- Barkasi, M., & Rosen, M. G. (2020). Is mental time travel real time travel? *Philosophy and the Mind Sciences*, 1, 1–27. <https://doi.org/10.33735/phimisci.2020.1.28>
- Barlow, H. B. (1956). Retinal noise and absolute threshold. *Journal of the Optical Society of America*, 46, 634–639. <https://doi.org/10.1364/JOSA.46.000634>
- Barlow, R. B., Birge, R. R., Kaplan, E., & Tallent, J. R. (1993). On the molecular origin of photoreceptor noise. *Nature*, 366, 64–66. <https://doi.org/10.1038/366064a0>
- Baylor, D. A., Matthews, G., & Yau, K. W. (1980). Two components of electrical dark noise in toad retinal rod outer segments. *The Journal of Physiology*, 309, 591–621. <https://doi.org/10.1113/jphysiol.1980.sp013529>
- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.

⁴ Many thanks to Lance Balthazar, Michael Barkasi, Sam Clarke, Kevin Lande, Heather Logue, Casey O’Callaghan, and two anonymous referees for written feedback, and to Andrew Buzzell, Brian Scholl, Scott Sturgeon, and especially Farid Masrouf for helpful conversations. This paper was discussed at the New Waves in Relationalism Conference in May 2021, where Farid commented. Although a full consideration of his rich and varied objections will have to wait for another occasion, I have in places attempted to clarify my position and arguments in reply to his comments. I’m grateful to him and many other conference participants whose comments helped me while revising this paper, including Rami Ali, Louise Antony, Michael Barkasi, Craig French, Ian Phillips, Thomas Raleigh, Kranti Saran, Susanna Schellenberg, and Wayne Wu. For helpful feedback in November 2021, I’m also grateful to Derek Brown, Fiona Macpherson, David Simmons, and other members of the University of Glasgow’s Philosophy, Psychology, Neuroscience Research Seminar. Portions of Section 2 of this article were recycled from an earlier, ill-conceived paper that I ultimately abandoned. Thanks to Wayne Wu, Alex Grzankowski, Steven Gross, two anonymous referees, and especially Peter Epstein and Kevin Lande for comments that helped me to eventually see that it was ill-conceived. My work on this paper was supported by Social Sciences and Humanities Research Council of Canada Insight Grant 435-2017-0923, and the Vision: Science to Applications Program, thanks in part to funding from the Canada First Research Excellence Fund.

- Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., & Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*, *13*, 87–100. <https://doi.org/10.1017/S095252380000715X>
- Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, *33*, 1–78. <https://doi.org/10.5840/philtopics20053311>
- Burge, T. (2010). *Origins of Objectivity*. New York: Oxford University Press.
- Byrne, A., & Logue, H. (2008). Either/or. In (A. Haddock & F. Macpherson Eds.), *Disjunctivism: Perception, Action, Knowledge* (pp. 57–94). Oxford: Oxford University Press.
- Byrne, A., & Manzotti, R. (draft). Hallucination and its objects. Unpublished manuscript.
- Campbell, J. (2002). *Reference and Consciousness*. New York: Oxford University Press.
- Chalmers, D. (2000). What is a neural correlate of consciousness? In (T. Metzinger Ed.), *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (pp. 17–40). Cambridge, MA: MIT Press.
- Chalmers, D. (2005). The matrix as metaphysics. In (C. Grau Ed.), *Philosophers Explore the Matrix* (pp. 132–176). Oxford: Oxford University Press.
- Cho, R., & Wu, W. (2013). Mechanisms of auditory verbal hallucination in schizophrenia. *Frontiers in Psychiatry*, *4*, 155. <https://doi.org/10.3389/fpsy.2013.00155>
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- Drouin, M., Kaiser, D. H., & Miller, D. A. (2012). Phantom vibrations among undergraduates: Prevalence and associated psychological characteristics. *Computers in Human Behavior*, *28*, 1490–1496. <https://doi.org/10.1016/j.chb.2012.03.013>
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429–433. <https://doi.org/10.1038/415429a>
- Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, *9*, 292–303. <https://doi.org/10.1038/nrn2258>
- ffytche, D. H. (2005). Visual hallucinations and the Charles Bonnet syndrome. *Current Psychiatry Reports*, *7*, 168–179. <https://doi.org/10.1007/s11920-005-0050-3>
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science*, *14*(5), 505–509. <https://doi.org/10.1111/1467-9280.03452>
- Green, D. M., & Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*. New York: John Wiley & Sons.
- Hellie, B. (2005). Noise and perceptual indiscriminability. *Mind*, *114*, 481–508. <https://doi.org/10.1093/mind/fzi481>
- Hellie, B. (2007). Factive phenomenal characters. *Philosophical Perspectives*, *21*, 259–306. <https://doi.org/10.1111/j.1520-8583.2007.00128.x>
- Hinton, J. (1967). Visual experiences. *Mind*, *76*, 217–227. <https://doi.org/10.1093/mind/LXXVI.302.217>
- Hoyer, P. O., & Hyvärinen, A. (2003). Interpreting neural response variability as Monte Carlo sampling of the posterior. *Advances in Neural Information Processing Systems*, *15* (*Proceedings of the 2002 International Conference on Neural Information Processing Systems*), 293–300.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*, 106–154. <https://doi.org/10.1113/jphysiol.1962.sp006837>
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, *55*, 271–304. <https://doi.org/10.1146/annurev.psych.55.090902.142005>
- Kriegel, U. (2013). Two notions of mental representation. In (U. Kriegel Ed.), *Current Controversies in Philosophy of Mind* (pp. 161–179). New York: Routledge.
- Langguth, B., Kreuzer, P. M., Kleinjung, T., & De Ridder, D. (2013). Tinnitus: Causes and clinical management. *The Lancet Neurology*, *12*, 920–930. [https://doi.org/10.1016/S1474-4422\(13\)70160-1](https://doi.org/10.1016/S1474-4422(13)70160-1)
- Langsam, H. (1997). The theory of appearing defended. *Philosophical Studies*, *87*, 33–59. <https://doi.org/10.1023/A:1002994402591>
- Lewis, D. (1980). Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy*, *58*, 239–249. <https://doi.org/10.1080/00048408012341251>
- Logue, H. (2012). Why naive realism? *Proceedings of the Aristotelian Society*, *112*, 211–237. <https://doi.org/10.1111/j.1467-9264.2012.00332.x>
- London, M., Roth, A., Beeren, L., Häusser, M., & Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, *466*, 123–127. <https://doi.org/10.1038/nature09086>

- Martin, M. G. F. (2002). The transparency of experience. *Mind & Language*, 17, 376–425. <https://doi.org/10.1111/1468-0017.00205>
- Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, 120, 37–89. <https://doi.org/10.1023/B:PHIL.0000033751.66949.97>
- Martin, M. G. F. (2006). On being alienated. In T. S. Gendler and J. Hawthorne (Eds), *Perceptual Experience* (pp. 354–410). Oxford: Oxford University Press.
- Masrour, F. (2020). On the possibility of hallucinations. *Mind*, 129, 737–768. <https://doi.org/10.1093/mind/fzy088>
- McDonnell, M. D., & Ward, L. M. (2011). The benefits of noise in neural systems: Bridging theory and experiment. *Nature Reviews Neuroscience*, 12, 415–425. <https://doi.org/10.1038/nrn3061>
- McDowell, J. (1982). Criteria, defeasibility and knowledge. *Proceedings of the British Academy*, 68, 455–79.
- Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Neander, K. (2017). *A Mark of the Mental*. In *Defense of Informational Teleosemantics*. Cambridge, MA: MIT Press.
- Noë, A. (2004). *Action in Perception*. Cambridge, MA: MIT Press.
- Noë, A. (2007). Magic realism and the limits of intelligibility: What makes us conscious. *Philosophical Perspectives*, 21, 457–474.
- O’Callaghan, C. (2007). *Sounds: A Philosophical Theory*. Oxford: Oxford University Press.
- Pajani, A., Kok, P., Kouider, S., & de Lange, F. P. (2015). Spontaneous activity patterns in primary visual cortex predispose to visual hallucinations. *Journal of Neuroscience*, 35, 12947–12953. <https://doi.org/10.1523/JNEUROSCI.1520-15.2015>
- Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.
- Papineau, D. (2014). Sensory experience and representational properties. *Proceedings of the Aristotelian Society*, 114, 1–33. <https://doi.org/10.1111/j.1467-9264.2014.00362.x>
- Pautz, A. (2010). Why explain visual experience in terms of content? In (B. Nanay Ed.), *Perceiving the World* (pp. 254–310). Oxford: Oxford University Press.
- Pautz, A. (2011). Can disjunctivists explain our access to the sensible world? *Philosophical Issues*, 21, 384–433. <https://doi.org/10.1111/j.1533-6077.2011.00209.x>
- Pautz, A. (2017). Experiences are representations: An empirical argument. In (B. Nanay Ed.), *Current Controversies in Philosophy of Perception* (pp. 23–42). New York: Routledge.
- Pautz, A. (2021). *Perception*. New York: Routledge.
- Penfield, W., & Perot, P. (1963). The brain’s record of auditory and visual experience. *Brain*, 86, 595–696. <https://doi.org/10.1093/brain/86.4.595>
- Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews: Neuroscience*, 1, 125–132. <https://doi.org/10.1038/35039062>
- Powers, A. R., Mathys, C., & Corlett, P. R. (2017). Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors. *Science*, 357, 596–600. <https://doi.org/10.1126/science.aan3458>
- Raleigh, T. (2014). A new approach to ‘perfect’ hallucinations. *Journal of Consciousness Studies*, 21, 81–109.
- Robinson, H. (1994). *Perception*. New York: Routledge.
- Rusakov, D. A., Savtchenko, L. P., & Latham, P. E. (2020). Noisy synaptic conductance: Bug or a feature? *Trends in Neurosciences*, 43, 363–372. <https://doi.org/10.1016/j.tins.2020.03.009>
- Schellenberg, S. (2010). The particularity and phenomenology of perceptual experience. *Philosophical Studies*, 149, 19–48. <https://doi.org/10.1007/s11098-010-9540-1>
- Schellenberg, S. (2011). Perceptual content defended. *Noûs*, 45, 714–750. <https://doi.org/10.1111/j.1468-0068.2010.00791.x>
- Schellenberg, S. (2018). *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18, 3870–3896. <https://doi.org/10.1523/JNEUROSCI.18-10-03870.1998>
- Shea, N. (2018). *Representation in Cognitive Science*. New York: Oxford University Press.
- Siegel, S. (2004). Indiscernability and the phenomenal. *Philosophical Studies*, 120, 91–112. <https://doi.org/10.1023/B:PHIL.0000033752.70521.13>
- Siegel, S. (2008). The epistemic conception of hallucination. In (F. Macpherson & A. Haddock Eds.), *Disjunctivism: Perception, Action, Knowledge* (pp. 205–224). Oxford: Oxford University Press.

- Snowdon, P. (1980). Perception, vision, and causation. *Proceedings of the Aristotelian Society*, 81, 175–92. <https://doi.org/10.1093/aristotelian/81.1.175>
- Soteriou, M. (2016). *Disjunctivism*. New York: Routledge.
- Sturgeon, S. (1998). Visual experience. *Proceedings of the Aristotelian Society*, 98, 179–200. <https://doi.org/10.1111/1467-9264.00031>
- Tinsley, J. N., Molodtsov, M. I., Prevedel, R., Wartmann, D., Espigulé-Pons, J., Lauwers, M., & Vaziri, A. (2016). Direct detection of a single photon by humans. *Nature Communications*, 7, 12172. <https://doi.org/10.1038/ncomms12172>
- Tomko, G. J., & Crapper, D. R. (1974). Neuronal variability: Non-stationary responses to identical visual stimuli. *Brain Research*, 79, 405–418. [https://doi.org/10.1016/0006-8993\(74\)90438-7](https://doi.org/10.1016/0006-8993(74)90438-7)
- Wiesenfeld, K., & Moss, F. (1995). Stochastic resonance and the benefits of noise: From ice ages to crayfish and SQUIDS. *Nature*, 373, 33–36. <https://doi.org/10.1038/373033a0>
- Wixted, J. T. (2020). The forgotten history of signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46, 201–233. <https://doi.org/10.1037/xlm0000732>
- Wolfe, J. M., Kluender, K. R., & Levi, D. M. (2015). *Sensation and Perception*, 4th Edition. Sunderland, MA: Sinauer Associates.

How to cite this article: Beck J. Mundane hallucinations and new wave relationalism. *Noûs*. 2021;1–23. <https://doi.org/10.1111/nous.12406>