How to Catch a Fade:
Biological Naturalism and the Fading Qualia Argument

Richard Brown

Abstract

Is consciousness a biological phenomenon, or an organizational invariant that any suitably organized system, a digital computer included, would share? The best argument I know for the second answer, computational functionalism, is Chalmers's fading qualia argument. If the biological answer were right, a rational, attentive subject with no functional pathology could be systematically wrong about its own experience. Chalmers says that the only tenable reply is to bite the bullet, and I argue that the bullet is smaller than it looks. He rests the premise that no such subject could be so wrong on the cases with which we are familiar, and what those cases give us reason to believe has changed with the evidence. Work on change blindness led Chalmers to stop treating his second argument, from dancing qualia, as a reductio, and that was his argument that a duplicate's experience would be like ours. Attentive observers miss slow changes in the very object they are asked to attend to, and fading is a slow change; these results, with others on subjective inflation and partial report, leave it open whether what observers say they see can come apart from how things look to them. Block has argued that the premise is either theory-laden or a generalization from cases that are all biological; I argue that it is now an open question whether those cases bear it out. No experiment shows an attentive subject misreporting a vivid, present experience, so a bullet remains. But the experiments have moved the bar for what counts as absurd here. None of this is evidence that consciousness is biological. It shows that the best argument against biological naturalism leaves it a live option.

1. Two Questions About Consciousness

Much energy in the philosophy of mind has been devoted to the question of whether consciousness is, or even could be, physical. There is an equally important question that is orthogonal to that one, and large language models have made it hard to avoid. Is consciousness fundamentally a biological phenomenon, or is it an 'organizational invariant' (Chalmers 1995), a property that any system with the right organization would share, whatever its substrate (Bostrom 2003)? On the second view, *computational functionalism*, a suitably programmed digital computer would be conscious in just the way that you or I ordinarily are. On the first, *biological naturalism*, some property of living systems is necessary, though perhaps not sufficient, for phenomenal consciousness (Seth 2026b, R1), and even a maximally specific computer simulation of the relevant biological processes would not thereby be conscious.

The two questions come apart (Chalmers 1995, §1); dualists and their opponents both divide over mine. In Block's (2015) terms it is orthogonal to the ontological question, whether there is anything immaterial. On his metaphysical question, what grounds experience, functionalism is a rival of physicalism, and there mine is a question about physicalism.

Biological naturalism has its defenders (Searle 2004; Block 2009, 2026; Seth 2026a), but computational functionalism is the more popular view, and it is argued for less often than one would expect. It is common to hear that neurons aren't magic, which is not an argument. The argument from multiple realizability at most suggests that there may be other *biological* configurations that are conscious. Michel (2026) puts the inductive case well. Functionalist theories explain the differences between types of mental state, and between conscious and unconscious states, and no biological theory does, so "as things stand we should bet on functionalism." That earns functionalism a share of one's credence, and it has a share of mine. But every one of those successes is with a living brain, and what explains a difference with the biology held fixed cannot show whether the biology is needed (Brown 2026, §V.c).

Nor will conceivability settle it. We can conceive of a conscious machine, a Commander Data. But conceiving shows only what is not ruled out by what we know. For a speaker ignorant of chemistry it is epistemically possible that water is XYZ (Putnam 1975), but given that water is H_2O, nothing else could have been water. If consciousness is a biological kind, nothing built as Data is could be conscious, however easily we conceive it. Conceiving of a machine with no experience at all gets the biological naturalist no further, and Chalmers himself says that the debate over such cases "tends to produce a stand-off" (1995, §2). He offers the fading qualia argument to break it.[1]

2. The Fading Qualia Argument

The argument starts with a fully biological and fully conscious creature. Call him Dave, and imagine that he has a migraine while watching a movie and eating a box of Sour Patch Kids. We replace one of his neurons with a silicon chip that "performs precisely the same local function as the neuron." It has "tiny transducers that take in electrical signals and chemical ions," computes on a digital signal, and converts the result into "the appropriate electrical and chemical outputs" (Chalmers 1995, §3). Then we replace two neurons, then three, and so on, the wonders of nanotechnology being what they are, until "there are no biochemical mechanisms playing an essential role" (1995, 314). Call the end of the series Robot and a case in the middle Joe. Each is

[1] This paper replaces an earlier one, "Consciousness Is (Probably) a Biological Phenomenon," which came out of discussion at the Fifth Online Consciousness Conference in 2013, and was presented in Taipei, and which I posted in 2019 (Brown 2019). Section 3 keeps its reading of Chalmers's change of mind, and Section 4 revisits its partial-report argument. I have changed the title because I am no longer as sure as its parenthesis suggests; I split my credence between a biological and a functional reading of any theory of consciousness (Brown 2025, 46, 48).

functionally isomorphic to Dave by assumption, so behavior cannot change. Joe goes on complaining about the headache and remarks that the sour candy takes his mind off the pain.

Suppose computational functionalism is false and Robot is not conscious. Then consciousness has to go away somewhere along the way. Chalmers and Block agree that it does not go all at once with the replacement of a single neuron (Chalmers 1995, §3; Block 2019, 573; 2023, 453). So it fades. Joe sees tepid pink where Dave sees bright red, and all the while he says and judges exactly what Dave does. He is "systematically wrong about everything that he is experiencing" (1995, §3). Chalmers allows that "in the extreme case, when all is dark inside," a system could be so wrong, since "in a sense there is nobody in there to be wrong" (1995, 316). The intermediate case is different.

> In every case with which we are familiar, conscious beings are generally capable of forming accurate judgments about their experience, in the absence of distraction and irrationality. For a sentient, rational being that is suffering from no functional pathology to be so systematically out of touch with its experiences would imply a strong dissociation between consciousness and cognition. We have little reason to believe that consciousness is such an ill-behaved phenomenon, and good reason to believe otherwise. (1995, 316)

Call this the *coherence premise*. Joe is rational, attentive, and free of functional pathology, so his judgments are accurate, his experience has not faded, and Robot is conscious after all. The argument is empirical; Chalmers grants that fading qualia may be logically possible, and argues from the cases we know (Block 2019, 574; 2023, 453–54). He says that his arguments "take the form of a reductio," but what they derive from the biological view is an implausibility and not a contradiction. They "do not constitute conclusive proof"; they are "strong plausibility arguments," and an opponent can hold on to absent qualia "only at significant cost" (1995, §2). He concludes that "the only tenable way for an opponent of organizational invariance to respond to this argument is to bite the bullet and accept the possibility of Fading Qualia" (1995, §3). I think he is right, unless experience goes all at once, an option that has its defenders (Mogensen 2025, §5) and that I set aside here.[2] The rest of this paper is about the size of the bullet.

Block replies that Joe is free of functional pathology only in a sense "that ignores phenomenal character," and that to prefer that sense is, as far as he can see, "to appeal to functionalism in a defense of functionalism" (2019, 575; 2023, 455). He adds that if something in "our biological makeup" is what causes our reports in "every case with which we are familiar," a case built so that something else causes them cannot refute that account of us (2019, 576–77; 2023, 456–57). I think that is correct. I would add, taking up a question that Block raises and puts aside (2019, 574, 576; 2023, 453–54, 457), that if each chip does everything the neuron did in its medium, including what it did to the ions and fields around it, the series never arrives at a digital

[2] I take it up in Brown (2026, n. 8).

computer, and that if it reproduces only the neuron's signals to other neurons, whether Joe is free of functional pathology is the point in dispute (Brown 2026, §VII; Brown manuscript). But my subject here is the evidence. Chalmers rests the premise on what the familiar cases give us reason to believe, and that is not what it was in 1995.

3. Flip-Flopping on Dancing and Fading Qualia

Chalmers has a second argument, and the two have different jobs. The fading argument, if it succeeds, shows that a functional isomorph of a conscious system is conscious, "but it does not establish that isomorphs have the same sort of conscious experience" (1995, §4). The second argument is meant to show that the experience is the same. A silicon backup for one of Joe's circuits is switched in and out, so that if silicon made any difference to experience, his experience would change with each flip and he could not notice. In 1995 Chalmers called that "a reductio ad absurdum." "As an empirical hypothesis," he wrote, "it is far more plausible that when one's experiences change significantly, then as long as one is rational and paying attention, one should be able to notice the change" (1995, §4). He now finds dancing qualia "very odd" but "something less than a reductio," since "work on change blindness has gotten us used to the idea that large changes in consciousness can go unnoticed." He adds that "these changes are outside attention, and unnoticed changes in the contents of attention would be much stranger," though "perhaps not so strange as to be ruled out in all circumstances," and he now takes the fading argument to be the stronger of the two (2010, 24 n. 7).

So Chalmers no longer treats it as absurd that Joe or Robot should have experience unlike Dave's with nothing they say or notice to show it, and the fading argument shows at most that they would have some experience. But the second argument was his better case for that as well. If Robot had no experience, he argued, a backup circuit could make experience dance "from vivid to tepid and back, without my ever noticing any change," and while an opponent "might bite the bullet and accept the possibility of Fading Qualia," dancing qualia were "an order of magnitude more difficult to accept" (1995, §4). The argument he now ranks first is the one whose bullet he thought an opponent might bite. What weakened the other was a set of experiments, and the fading argument's central premise is the same kind of claim about what is unprecedented.

The experiments go further than Chalmers allows. When a picture is changed during a blink, observers told to expect a change miss it more than 40 percent of the time even when their eyes are on the spot where it happens (O'Regan et al. 2000, 200). When the change is made during an eye movement, observers who are studying the picture for a memory test, and have been told to press a button at any change, still miss it 42 percent of the time when their eyes have rested on the changed thing both before and after. Asked, once the picture is gone, which of four statements was true of it, they choose the feature it had before the change on a third of those

misses, about as often as the feature it had after (Lee et al. 2025, 23, 25–26). Two of the study's authors read these misses in opposite ways. Rosenthal, for the higher-order theory, finds it "overwhelmingly likely" that some who chose the old feature "were conscious of perceiving a feature different from the feature they actually perceived"; Haun, for integrated information theory, holds that a change goes unnoticed for want of the states that comparison takes, "rather than because the change was entirely absent from experience" (Lee et al. 2025, 52, 47). The usual explanation of the misses is that one can look at a thing, and even attend to it, without taking in the feature that changes (O'Regan et al. 2000, 204; Simons and Levin 1998, 648), so Chalmers's qualification survives only for the very feature one is attending to.

Slow change presses even on that, and it bears on fading directly, because fading is a slow change. If the series is one person's story, as I told it above, Joe's experience dims by degrees while he attends to it, and he notices nothing, which is what Chalmers's hypothesis says should not happen. The gradual changes in one early study were thought visible to anyone attending to the right feature or region (Simons, Franconeri, and Reimer 2000, 1152). But when a large region at the center of a scene, a car or a wall, changed color over sixteen seconds, all but 14 of 417 observers, told that features of the image might change, reported no change in it (Frey et al. 2024, 10). And observers told to pay attention to a couch in the middle of a scene, though not to its color, mostly failed to notice its hue drifting by as much as 60 degrees; the authors think that attention made them more blind to the change, not less (Frey and Whitney 2026).

What became of these observers' experience is not settled. Probed just after the scene had gone blank, those who had missed the change accepted the region's last color as a match more often than its first, and on the authors' account they had formed "a rich but fragile and fleeting representation of the object as it just was on the screen" beside a coarse and stable one (Frey et al. 2024, 13–14). Those who missed the couch's drift, asked its color while it was in front of them and choosing from a wheel once it was gone, picked a hue that had moved less than a fifth as far as the couch had (Frey and Whitney 2026). If how the object looked went with the fleeting representation, their experience changed by degrees before their eyes and they did not notice, which is how Chalmers describes change blindness and what his hypothesis rules out. If it looked the way they went on to report it, pulled toward the past, which is Frey and Whitney's reading and mine, they were wrong about its color and right about their experience. The experiments do not say which. On either reading the fading case is harder. These observers are wrong, at most, about whether anything changed, or about a color just after it was in view, and Joe is wrong about what he is experiencing now, a bright red and a migraine, with as long as he likes to look.

Nor does the general case against introspection reach him. Schwitzgebel (2010) set change blindness against the dancing argument, at least if experience is rich, granting Chalmers's point that many cases need attention to be off the thing that changes and replying that not all seem to;

against the fading argument he set the dispute over how rich experience is, in which one side or both must be radically wrong about their own experience. Yet Schwitzgebel, who argues that "we're prone to gross error, even in favorable circumstances of extended reflection," grants that it is "hard, seemingly, to go too badly wrong in introspecting really vivid, canonical pains and foveal colors" (2008, 259–60), and those are Joe's cases. So a reply cannot rest on evidence of misreport. His own later reply does not. It is that the fading argument gives a reader who doubts Robot's reports no reason to trust Joe's (Schwitzgebel 2022, 2024, adapting Udell and Schwitzgebel 2021). Mine rests on what these and two other sets of results leave open.

4. What Inflation and Partial Report Leave Open

In a recent preregistered adversarial collaboration, whose authors include Block, Chalmers, and me, observers judged the orientation of briefly presented peripheral targets and reported how visible they were (Tian et al. 2025). Across the same stimuli, performance fell from 87 to 66 percent correct when attention was directed elsewhere, while reports of how visible the stimulus was changed far less. At matched levels of performance, unattended stimuli were reported as more visible than attended ones, in all four experiments, and above threshold as well as near it.

This does not show that anyone misreports their experience. For human beings I take reports at face value, and that includes the surprising ones; that is how Lau and I read earlier results of this kind (Lau and Brown 2019; Brown 2025, 220). So the dissociation is between how visible things are reported to be and how well they are discriminated. The paper leaves open what the reports go with instead, and says that the result "is not on its own necessarily incompatible with first-order theories" (Tian et al. 2025). On a higher-order account the report runs through a higher-order state, and how things look is settled there. The results favor that account only if no first-order activity carries what subjects report and the inflation is in how things look, not just in the ratings, and neither has been shown (Brown 2026, n. 1). They are also neutral between a functional and a biological reading of the account, since the two predict the same reports. On the biological reading, which I explore in a companion paper and do not argue for here, a higher-order representation is phenomenally conscious only when it has the right biological realization (Brown 2026, §§V.c–d).

The inflation in these experiments is an effect of inattention, and Joe is attending. But inflation does not take inattention. Observers are more willing to say that a target is present in the periphery than at the center of gaze, even when the target in the periphery is one they have been asked to detect (Solovey, Graney, and Lau 2015; Odegaard et al. 2018).

So on the reading I prefer nobody in these experiments is wrong about their experience, and the results are no precedent for Joe. But that reading is not forced. If no first-order activity carries what these observers report, there is a mismatch either way. Put the experience where the report

comes from, and it outruns perception, which is inflation; the higher-order theorist serves that bullet as the main course. Put it in the perceptual state, and it falls short of what these observers take themselves to experience, which is fading, Chalmers's bullet on a small scale. Whether inflation is in how things look or in how observers decide what to say is in dispute (Abid 2019; Knotts, Michel, and Odegaard 2020), and so is whether the pull toward the recent past arises in perception, where a trace of it has been found as early as primary visual cortex, or in the decision that follows it (St. John-Saaltink et al. 2016; Manassi and Whitney 2022; Fritsche, Mostert, and de Lange 2017). On the second answer to each, what these observers say they see comes apart from how things look to them. They are rational and free of functional pathology, and with the couch and the targets they were asked to detect, they are attending to the object in question. That would be a small case of what the coherence premise says we have little reason to expect.

Partial report raises the same question about everything outside the focus of attention. When a large set of letters is flashed, "we feel that we can see them all, while we can report only a few," and observers go on feeling that there were only letters when rotated or flipped ones are put in the rows they are not asked about (de Gardelle, Sackur, and Kouider 2009, 569). Kouider and his colleagues conclude that the richness is an illusion built from fragments and expectation (Kouider et al. 2010), and Block, for whom "the illusion effect is small," grants that there is one (2011, 569–70). It is not only brief displays. When color was slowly removed from everything outside a small region around the point of gaze, close to a third of observers exploring a scene noticed nothing, even with less than 5 percent of the display in color (Cohen, Botch, and Robertson 2020). Set that beside inflation. Either the periphery went on looking colored with no color in it, or it did not, and they did not know it. If it did not, then since they were looking around, each thing's color faded as their eyes moved off it and came back when they returned to it, and they noticed none of it. If the periphery is ordinarily like that, experience fades in all of us whenever we look around, and we do not catch it. I once took partial report to show that normal subjects can have partly degraded experiences without knowing it (Brown 2019). It shows that only if how things look goes with the fragments. On the rich version of the higher-order account, these observers experience letters, as they say (Brown 2025, 15, 107). On the sparse version, outside the focus of attention we "experience some general or partial content but experience it as though it presented a great amount of detail," and the difference from the rich case may be "subjectively indistinguishable" (2025, 105, 107 n. 10). Less would then be phenomenologically manifest to us than we would say, and we might have no way to tell. That is Joe's condition in miniature. And the sparse reading does not take a higher-order theory; Kouider and his colleagues hold a version of it on an account of access. My sympathies are with the rich version, but we cannot yet decide between the two.

When I first put this argument to Chalmers, he replied that fading "requires huge ongoing errors," and that, where the switching explains the momentary errors of the dancing case, here "there's no analogous explanation of those errors, in that the system will be functionally innocuous throughout" (Chalmers 2013). It is true that each of the mismatches above has a functional cause that a duplicate would share, and that Joe's has none. That is why none of them is a counterexample. It does not make Joe's error inexplicable, since the series is built to preserve whatever his judgments depend on (cf. Block 2019, 577; 2023, 457), and unreliability "would simply be what we should expect" (Schwitzgebel 2024; cf. 2010). What the mismatches bear on is how far the sense of absurdity can be trusted. All I claim is that the experiments have moved the bar. In 1995 it could be said that in every familiar case judgment generally keeps step with experience, distraction, irrationality, and pathology aside. That is now one of two live readings of what happens in the laboratory to attentive subjects with whom nothing has gone wrong. Each result has part of Joe's condition and none has all of it, and how much smaller the bullet is depends on the credence the other reading deserves. But if something like Joe's condition may be going on in us, it is less absurd to think it is going on in him. What looked like a tooth-shattering bullet turns out to be a step closer to an unusual amuse-bouche.

5. Conclusion

The fading and dancing qualia arguments are plausibility arguments. What they derive from biological naturalism is an improbable dissociation between experience and judgment about experience, and how improbable it is depends on the evidence. I have argued that the evidence has moved, once on Chalmers's own account and again since. No experiment shows an attending subject who is wrong about a vivid, present experience, so there is still a bullet to bite. But it is now an open question whether small versions of Joe's condition occur in us, and what is possibly true of us, given what we know, cannot be so absurd in him. Chalmers says, of the intuition that some functional duplicate of a brain would not be conscious, that intuition "is unreliable as a guide to empirical possibility, especially where a phenomenon as perplexing as conscious experience is concerned" (1995, §2). The history of his second argument suggests that the same goes for the intuition that Joe could not be so wrong. None of this is evidence that biological naturalism is true. I still split my credence between a biological and a functional reading of the higher-order theory (Brown 2026, §VIII). But the best argument against biological naturalism leaves it a live option.

What follows for artificial intelligence is the subject of the companion paper (Brown 2026), where I argue that if biological naturalism is true, Robot would be no mere imitator of a mind; it would have the kind of consciousness that illusionists think we have, real and self-representing but not phenomenal. Science has shown us that the world is far stranger than any of us could

have ever imagined. Whether it is strange in the way the biological naturalist requires is an empirical question.[3]

References

Abid, G. (2019). Deflating inflation: The connection (or lack thereof) between decisional and metacognitive processes and visual phenomenology. *Neuroscience of Consciousness* 5(1): niz015. doi:10.1093/nc/niz015.

Block, N. (2009). Comparing the major theories of consciousness. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences* (4th ed., pp. 1111–1122). MIT Press.

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences* 15(12): 567–575.

Block, N. (2015). The Canberra Plan neglects ground. In T. Horgan, M. Sabatés, and D. Sosa (eds.), *Qualia and Mental Causation in a Physical World: Themes from the Philosophy of Jaegwon Kim*, 105–133. Cambridge University Press.

Block, N. (2019). Fading qualia: A response to Michael Tye. In A. Pautz and D. Stoljar (eds.), *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness*, 571–580. MIT Press.

Block, N. (2023). *The Border Between Seeing and Thinking*. Oxford University Press.

Block, N. (2026). Can only meat machines be conscious? *Trends in Cognitive Sciences* 30(4): 298–308.

Bostrom, N. (2003). Are you living in a computer simulation? *Philosophical Quarterly* 53(211): 243–255.

Brown, R. (2019, June 19). Consciousness is (probably) a biological phenomenon. Blog post. https://onemorebrown.com/2019/06/19/consciousness-is-probably-a-biological-phenomenon/

Brown, R. (2025). *Consciousness as Representing One's Mind: The Higher-Order Approach to Consciousness Explained*. New York: Oxford University Press.

Brown, R. (2026). Real consciousness or real illusion? A higher-order approach to AI consciousness. Manuscript. https://philarchive.org/rec/BRORCO-5

Brown, R. (manuscript). Biologically based: Brain rhythms and biological theories of consciousness.

Chalmers, D. J. (1995). Absent qualia, fading qualia, dancing qualia. In T. Metzinger (ed.), *Conscious Experience*, 309–328. Paderborn: Ferdinand Schöningh.

Chalmers, D. J. (2010). *The Character of Consciousness*. Oxford University Press.

Chalmers, D. J. (2013, February 19). Comment on M. Á. Sebastián, 'What panpsychists should reject: On the incompatibility of panpsychism and organizational invariantism.' Fifth Online Consciousness Conference. https://consciousnessonline.wordpress.com/2013/02/15/what-panpsychists-should-reject-on-the-incompatibility-of-panpsychism-and-organizational-invariantism/

Cohen, M. A., Botch, T. L., & Robertson, C. E. (2020). The limits of color awareness during active, real-world vision. *Proceedings of the National Academy of Sciences* 117(24): 13821–13827. doi:10.1073/pnas.1922294117.

de Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition* 18(3): 569–577.

Frey, H. G., Koenig, L., Block, N., He, B. J., & Brascamp, J. W. (2024). Memory representations during slow change blindness. *Journal of Vision* 24(9): 8, 1–18. doi:10.1167/jov.24.9.8.

[3] This paper benefited from drafting, editing, suggestions, and discussion with Claude.

Frey, H. G., & Whitney, D. (2026). Slow change blindness from serial dependence: Slowly changing objects are perceived as biased toward the past. *BMC Biology* 24: 225. doi:10.1186/s12915-026-02700-w.

Fritsche, M., Mostert, P., & de Lange, F. P. (2017). Opposite effects of recent history on perception and decision. *Current Biology* 27: 590–595. doi:10.1016/j.cub.2017.01.006.

Knotts, J. D., Michel, M., & Odegaard, B. (2020). Defending subjective inflation: An inference to the best explanation. *Neuroscience of Consciousness* 6(1): niaa025. doi:10.1093/nc/niaa025.

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences* 14(7): 301–307.

Lau, H., & Brown, R. (2019). The Emperor's new phenomenology? The empirical case for conscious experience without first-order representations. In A. Pautz & D. Stoljar (Eds.), *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness* (pp. 171–198). MIT Press.

Lee, I., Sans, A., Ng, L., Faulkner, R., Haun, A., Chesney, D., Rosenthal, D., Fallon, F., Odegaard, B., & Lee, A. L. F. (2025). Change detection and saccades: Replicating and extending the work of John Grimes. *PsyArXiv*. doi:10.31234/osf.io/f368e_v1.

Manassi, M., & Whitney, D. (2022). Illusion of visual stability through active perceptual serial dependence. *Science Advances* 8(2): eabk2480.

Michel, M. (2026). Bet on functionalism. Commentary on Seth (2026a). *Behavioral and Brain Sciences* 49, e348. doi:10.1017/S0140525X25103816.

Mogensen, A. L. (2025). How to resist the fading qualia argument. *Synthese* 206: 252. doi:10.1007/s11229-025-05338-3.

Odegaard, B., Chang, M. Y., Lau, H., & Cheung, S.-H. (2018). Inflation versus filling-in: Why we feel we see more than we actually do in peripheral vision. *Philosophical Transactions of the Royal Society B* 373(1755): 20170345.

O'Regan, J. K., Deubel, H., Clark, J. J., & Rensink, R. A. (2000). Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition* 7: 191–211.

Putnam, H. (1975). The meaning of 'meaning.' *Minnesota Studies in the Philosophy of Science* 7: 131–193.

Schwitzgebel, E. (2008). The unreliability of naive introspection. *Philosophical Review* 117(2): 245–273.

Schwitzgebel, E. (2010, April 22). Chalmers's fading/dancing qualia and self-knowledge. Blog post, *The Splintered Mind*. https://schwitzsplinters.blogspot.com/2010/04/chalmerss-fadingdancing-qualia-and-self.html

Schwitzgebel, E. (2022, December 13). An objection to Chalmers's fading qualia argument. Blog post, *The Splintered Mind*. https://schwitzsplinters.blogspot.com/2022/12/an-objection-to-chalmerss-fading-qualia.html

Schwitzgebel, E. (2024, October 22). An objection to Chalmers's fading qualia argument. Substack post, *The Splintered Mind*. https://eschwitz.substack.com/p/an-objection-to-chalmerss-fading

Searle, J. R. (2004). *Mind: A Brief Introduction*. Oxford University Press.

Seth, A. K. (2026a). Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences* 49, e315 (with commentaries and response; first published online 2025). doi:10.1017/S0140525X25000032.

Seth, A. K. (2026b). The stuff matters: Consciousness, computation, and biology. Author's response to commentaries on Seth (2026a). *Behavioral and Brain Sciences* 49, e366. doi:10.1017/S0140525X26106906.

Simons, D. J., Franconeri, S. L., & Reimer, R. L. (2000). Change blindness in the absence of a visual disruption. *Perception* 29: 1143–1154.

Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review* 5: 644–649.

Solovey, G., Graney, G. G., & Lau, H. (2015). A decisional account of subjective inflation of visual perception at the periphery. *Attention, Perception, & Psychophysics* 77(1): 258–271.

St. John-Saaltink, E., Kok, P., Lau, H. C., & de Lange, F. P. (2016). Serial dependence in perceptual decisions is reflected in activity patterns in primary visual cortex. *Journal of Neuroscience* 36(23): 6186–6192.

Tian, K. J., Maniscalco, B., Epstein, M. L., Shen, A., Graham Castaneda, O., Kurosawa, T., Motzer, J. A., Olsson, E., Russell, E. E., Walsh, M. E., Wang, J., Bek Awrang Zeb, T., Brown, R., Lamme, V. A. F., Lau, H., He, B. J., Brascamp, J. W., Block, N., Chalmers, D., Peters, M. A. K., & Denison, R. N. (2025). When awareness outstrips performance: Critical tests of subjective inflation under inattention. *bioRxiv*. doi:10.1101/2025.07.03.661972 (version 2, January 24, 2026).

Udell, D. B., & Schwitzgebel, E. (2021). Susan Schneider's proposed tests for AI consciousness: Promising but flawed. *Journal of Consciousness Studies* 28(5–6): 121–144.