



On track to a standard model

Daniel C. Dennett

To cite this article: Daniel C. Dennett (2020) On track to a standard model, *Cognitive Neuropsychology*, 37:3-4, 173-175, DOI: [10.1080/02643294.2020.1731443](https://doi.org/10.1080/02643294.2020.1731443)

To link to this article: <https://doi.org/10.1080/02643294.2020.1731443>



Published online: 09 Mar 2020.



Submit your article to this journal [↗](#)



Article views: 243



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)

COMMENTARY



On track to a standard model

Daniel C. Dennett

Center for Cognitive Studies, Tufts University, Medford, MA USA

ARTICLE HISTORY Received 3 December 2019; Accepted 5 February 2020

KEYWORDS Consciousness; Hard Problem; Hard Question; first-person; illusionism

I applaud the ambition and optimism with which the authors lay out their case for a “single, coherent explanation of consciousness” that can unite GW, HOT and illusionism into a “standard model” based on their Attention Schema Theory (AST). I think they are basically right, and although I have some misgivings about a few of their ways of putting things, and think they underestimate the scope and complexity of their own theory, in the spirit of advancing their project, I will devote most of my commentary to adding what I take to be friendly amendments and clarifications.

I will highlight the themes I find most important. “I-consciousness is what the brain actually has; m-consciousness is what the brain thinks that it has.” (p. 13) Yes, i-consciousness is the dynamic informational system that is physically realized in the brain (in one variety of global workspace or another), and that accounts for all the behaviours, talents and dispositions that persuade us that an organism is conscious in one way or another. Chalmers’ “easy problems” (Chalmers, 1995) are all solved or solvable in a developed theory of i-consciousness. But it is somewhat misleading to say that *the brain* thinks that it has m-consciousness; only (the brains of) *people* – and particularly reflective people – think that they have m-consciousness. Dogs’ brains have i-consciousness systems similar in many regards to those in our brains – they have attention-systems controlled by competitive sub-assemblies that determine what gets enhanced when – but their brains don’t generate the theorists’ illusion that our brains do, primarily because dogs don’t ever learn to compare notes with their conspecifics, to communicate about what is going on in their brains.

Dogs presumably do not think there is something it is like to be them, even if there is. It is not that a dog thinks there isn’t anything it is like to be a dog; the dog is not a theorist at all, and hence does not suffer from the theorists’ illusion. The hard problem and meta-problem are only problems for us humans, and mainly just for those of us humans who are particularly reflective. In other words, dogs aren’t bothered or botherable by problem intuitions. Dogs – and, for that matter, clams and ticks and bacteria – do enjoy (or at any rate benefit from) a sort of user illusion: they are equipped to discriminate and track only some of the properties in their environment. (Dennett, 2019, p. 54)

The authors speak of the “cultural ubiquity” of the m-consciousness belief, which nicely draws attention to the proposition that m-consciousness is a *myth*, just as much an *artifact* of human sociality and cultural exchange as is religion or music. The fact that it is such a natural byproduct of human communication, based as it is on deep facts about the i-consciousness we humans share, leads most theorists to *intuit* (not infer) that it is simply undeniable. But as the authors note (p. 23), *all* the “arguments” for the existence of m-consciousness are patently circular.

The illusion of m-consciousness (Frankish, 2016, 2017) is a myth, but not a mere cultural ornament; it is “the brain’s quick-and-dirty, but useful model of i-consciousness.” (p. 14) This shared noticing of our own attention schemas enables us to control our attention in ways that dogs and members of other species cannot. Our human capacity for self-control (and the control of our conspecifics) dwarfs that of all other species and is the foundation for our culturally developed and endorsed traditions of moral responsibility. This unique talent depends on our having an attention-schema that models (sketchily)

the dynamic processes of attention-allocation, permitting us – and only us – to *notice* when the neural competitions for attention get out of balance, captured by some accidental or possibly malign intrusion, and then to act on our noticing, by *directing* our attention elsewhere. “When you are aware of something, you can choose to act on it.” (p. 20) Dogs and other animals do exhibit some modest capacities for noticing their noticings, but we humans have mental lives that teem with such episodes – so much so that most people have never even imagined that the mental lives of other species might not be similarly populated. Watching a raccoon figure out how to screw off a trashcan lid, for instance, we readily furnish its mind with hunches and surmises and hopes and expectations *that we suppose it consciously entertains*, without acknowledging so much as the possibility that the raccoon may get the benefits of an internal process of trial-and-error (it is a Popperian creature, trying things out in its head before acting, Dennett, 1995) *without realizing it*. As the authors note, “We have a hair trigger for attributing consciousness, because it is so socially useful that it is better to mistakenly overuse it than mistakenly underuse it.” (p. 30) One of the merits of AST is that it admits of grades and degrees. Instead of asking whether the lightbulb of consciousness is ON or OFF, it asks what levels of control-of-attention are implicated in the repertoire of this animal or person.

Why has the intimate link between consciousness and control not been more widely explored? Here’s a hypothesis: most theorists of consciousness don’t even try to ask what I call the Hard Question: *And then what happens?* (Dennett, 2018). They develop the “inbound” half of their models – getting from photons and pressure waves all the way up to *subjective experience* (whatever that is) and then they stop. They sometimes even *endorse* stopping, so as to isolate “the neural basis of conscious perception itself” from all “post-perceptual cognitive processing.” (Block, *forthcoming*) But you can’t understand the power of a key by studying it in isolation; you have to study the lock into which it is designed to fit.

Another reason why the link is ignored, I surmise, is because the prospect of investigating the underlying processes of control of attention seems to fly in the face of our misguided convictions about m-consciousness. Tradition has it that “we” are somehow “directly”, “incurably”, or even “infallibly” in touch with the

“phenomenal properties” of our subjective experience. We tend to think that no relation is more intimate than the relation between us and our “qualia”. But a scientific account (from the “third-person point of view”) of the underlying machinery, which *must be there* to get “us” from subjective experience of an apple to our ability to tell others – or ourselves – that we are experiencing an apple, must necessarily renounce all pre-susptions of magical directness. The authors note:

Like the body schema, the attention schema is constructed automatically. We do not have cognitive control over it. Also like the body schema, at least some aspects of the proposed attention schema are cognitively accessible and verbalizable. (p. 27)

Some aspects – content aspects – are “cognitively accessible”, and the rest *are not*. It is important to note that, as the AST makes clear, cognitive *inaccessibility* is always the default case. It takes extra machinery, an extra layer of representation-consumption and subsequent control, to make something accessible “to us”.

To see the awkwardness of asking the Hard Question, it helps to walk through a simple case from start to finish:

Spy (on his cell phone to his handler): Smith’s house has a dark blue door.
 Handler: How do you know?
 S: I’m looking at it right now, and there’s plenty of light, and my GPS tells me that this is indeed Smith’s address – oh, and my eyes are open.
 H: OK, so the light comes in your eyes. Then what happens?
 S: Well, I ... *see the blue door*. I know that this depends on a lot of activity in the retina, optic nerves, lateral geniculate nucleus, striate cortex, ... the temporoparietal junction, but that is all stuff I’ve learned about second-hand from reading cognitive science books and articles. All I know *directly* and *intimately* is that I see the door. Well, that’s too strong, I guess; I *seem to see the blue door*. My subjective experience is of a blue door.
 H: How do you know that is your subjective experience?
 S: What do you mean? I just do know!
 H: That won’t do, I’m afraid. You’ve just now vehemently expressed your opinion about your subjective experience, but I want to know how you were enabled to do that.

Here are two things S might then say:

M-C: I am directly acquainted with a phenomenal property or quale, *subjective blue*, and another, which might be called *door-shape straight ahead*. It is my direct acquaintance with these properties that enables me to tell you about them (of course).

I-C: Well, I've never seen any details on this, but presumably my brain was able to, um, discriminate some internal state as a representation of blue, and ... enable another part of my brain to activate the English word "blue" and endorse a link of some kind between the word and the representation (ruling out competing alternative links as somehow inferior) ... but it must be much more complicated than that *and none of the details are cognitively accessible to me*. I don't know how I am caused to have the conviction I just expressed, but I still have it.

S should abjure the M-C option, since it is nothing more than "a lingering fragment of a larger cluster of physically incorrect beliefs." (p. 8) An important point that should not be overlooked is that S and H have to learn how to do all this. We human beings aren't born knowing how to talk about, or direct our attention to, or discriminate ... our subjective experiences. There has to be a process that is both cognitive and social that sculpts our ability to "introspect" and even as adults we can discover embarrassing gaps of "ineffability" that can be filled in with training and practice. Acquiring all those talents installs cognitive machinery in our brains that is useful not just for sharing experiences with our family and friends, but for controlling our attention in myriad ways. (See Dor, 2015 for a rich and imaginative theory of this. See also Markkula, 2015.)

The authors actually address the dual task confronting any theorist of consciousness: putting the third-person, physical scientific model of i-consciousness in registration, somehow, with "I"-talk and "we"-talk. Some otherwise insightful theorists often simply punt:

We can't possibly know (let alone keep track of) the tremendous number of mechanical influences on our behavior because *we inhabit* an extraordinarily complicated machine. (Wegner, 2002, p. 27)

I never have the sort of direct access that my mindreading system has to *my own* visual images and bodily feelings. (Carruthers, 2009, sec.2 para.8) (For discussion, see Huebner & Dennett, 2009)

(Emphases added in both quotations. Where did these "I"s and "we"s come from?) Both Wegner and

Carruthers have made significant contributions to the clarification of the relationship between what I have called the personal and subpersonal levels, but they give themselves poetic license to make these entirely reasonable assertions without noting the background presumption that is needed to keep these observations from being dualist nonsense: the personal level is itself a user-illusion, one of nature's greatest inventions, and they are exploiting that very illusion in making these claims.

There is much more in this essay that deserves careful discussion, but first we should welcome it and take it seriously as a worthy candidate for a "standard model" of human consciousness, not just a bold philosophical conjecture but an empirically buttressed and detailed articulation of what should now be recognized as the "obvious default theory" (Dennett, 2016), incorporating many of the fruits of recent efforts across several fields, and well positioned to incorporate many more.

References

- Block, N. (forthcoming). *Trends in cognitive sciences*, "What is wrong with the no-report paradigm and how to fix it".
- Carruthers, P. (2009). How we know our own minds: the relationship between mindreading and metacognition.
- Chalmers, D. (1995). Facing up to the hard problem of consciousness. *Journal of Consciousness Studies*, 2, 200–219.
- Dennett, D. (1995). *Darwin's dangerous idea*. New York: Simon & Schuster.
- Dennett, D. (2016). Illusionism as the obvious default theory of consciousness. *Journal of Consciousness Studies*, 23(11–12), 65–72.
- Dennett, D. (2018, July 30). Facing up to the hard question of consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373, 20170342. doi:10.1098/rstb.2017.0342.
- Dennett, D. (2019). Welcome to strong illusionism. *Journal of Consciousness Studies*, 26(9–10), 48–58.
- Dor, D. (2015). *The instruction of imagination: language as a social communication technology*. Oxford: Oxford Univ. Press.
- Frankish, K. (2016). Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11–12), 11–39.
- Frankish, K. (Ed.). (2017). *Illusionism as a theory of consciousness*, *Journal of Consciousness Studies*, special issue, reprinted by Imprint Academic, PO Box 200, Exeter EX5 5YX UK.
- Huebner, B., & Dennett, D. (2009). Banishing 'I' and 'we' from accounts of metacognition. *Behavioral and Brain Sciences*, 32, 148–149. doi:10.1017/S0140525X09000661
- Markkula, G. (2015, June 16). Answering questions about consciousness by modeling perception as covert behavior. *Frontiers in Psychology*, 6. doi:10.3389/fpsyg.2015.00803.
- Wegner, D. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.