

- Lewis, David. 1979. "Counterfactual Dependence and Time's Arrow." *Noûs* 13, no. 4: 455–76.
- List, Christian, and Peter Menzies. 2009. "Nonreductive Physicalism and the Limits of the Exclusion Principle." *Journal of Philosophy* 106, no. 9: 475–502.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- Zhong, Lei. 2020a. "Intervention, Fixation, and Supervenient Causation." *Journal of Philosophy* 117, no. 6: 293–314.
- Zhong, Lei. 2020b. "Taking Emergentism Seriously." *Australasian Journal of Philosophy* 98, no. 1: 31–46.

Lei Zhong

The Chinese University of Hong Kong

Philosophical Review, Vol. 130, No. 4, 2021

DOI 10.1215/00318108-9264056

Peter Carruthers, *Human and Animal Minds: The Consciousness Questions Laid to Rest*.

Oxford: Oxford University Press, 2019. 220 pp.

Peter Carruthers has a new book. Yes, you guessed right: the title ends with the word *Mind*, and you'll probably disagree with the conclusion. *The Architecture of the Mind* concludes that the mind is massively modular (Carruthers 2006). *The Opacity of Mind* concludes that self-knowledge (of propositional attitudes) is unattainable (Carruthers 2011). *The Centered Mind* concludes that there's no such thing as conscious thought (Carruthers 2015). Now, in *Human and Animal Minds*, Carruthers concludes that there's no fact of the matter as to whether nonhuman animals are phenomenally conscious or not.

Even if you don't like the final destination, reading a book by Carruthers is always a beautiful journey, owing to the arguments and pieces of evidence gathered along the way. *Human and Animal Minds* is no exception.

On this journey you'll find Carruthers's reflections on the current evidence for metacognition in nonhuman animals (chapter 2); his rebuttal of 'theory-neutral' approaches to nonhuman animal consciousness (chapter 3); his arguments against theories of consciousness such as integrated information theory, local recurrence theory, and higher-order theories of consciousness—which Carruthers himself used to hold (chapter 4); and his defense of the view that consciousness doesn't really matter for ethics (chapter 8).

Perhaps most importantly, you'll get a chance to taste Carruthers's new reductionist cocktail, which threatens to explain phenomenal consciousness by

mixing two dangerous substances: the phenomenal concept strategy and the global workspace theory of consciousness (chapters 5 and 6).

Here's how Carruthers's version of global workspace theory works (I focus only on visual consciousness). You become phenomenally conscious of a visual feature when a nonconceptual representation of that feature is encoded into a "global workspace" which globally broadcasts this representation to a broad set of neurocognitive modules. This global broadcast mechanism is all-or-none: representations don't get to be half-globally broadcast.

Once a nonconceptual content is globally broadcast, it is available as a target for a higher-order demonstrative thought using a purely recognitional concept—a phenomenal concept—of the kind "*This reddish experience (This-R).*" Since phenomenal concepts are purely recognitional (nondescriptive) concepts, they are conceptually isolated from other concepts—for instance, concepts referring to brain states or psychological functions.

This conceptual isolation explains why you can think "I could have globally broadcast nonconceptual content without *This-R*," even if '*This-R*' in fact refers to globally broadcast nonconceptual content. It also explains why you believe in the explanatory gap, why you can conceive of zombies, and why Mary can't deduce what it's like to see red from her physical knowledge while in her black-and-white room. Phenomenal consciousness just *is* globally broadcast nonconceptual content. There's no gap in nature. Just a gap in understanding. And as long as there's no gap in nature, physicalists can rest easy in their knowledge that phenomenal consciousness doesn't involve any spooky properties.

To understand the next step in Carruthers's argument, I'm asking you to think about baldness. Some of my friends are entering this stage of capillary life where they're *kind of* bald. They're not *really* bald. Just *kind of* bald. If you don't have any kind-of-bald friends, you can still conceive of that property. In contrast, you can't conceive of a figure that'd be *kind of* a triangle. And you can't conceive of a number that'd be *kind of* nine. So in that respect there's a difference between the concepts *bald*, and *triangle*, or *nine*.

Is the concept *phenomenal consciousness* more like *bald*, or more like *triangle*? Carruthers says it's the latter: "we can't make sense of degrees of phenomenal consciousness" (23). What would it be like to *kind of* have a conscious experience, but *not really*?

Of course, I can conceive being conscious *of* a very weak stimulus, like experiencing just a dim light. But there's a difference between being conscious of a dim light and being, so to speak, *dimly* conscious of a light. We all agree that what you are conscious *of*—the *contents* of consciousness—can come in degrees. The question is whether phenomenal consciousness itself is graded. And *that* is difficult to imagine.

Carruthers concludes that *phenomenal consciousness* is an all-or-none concept. This fits well with global workspace theory and its all-or-none global broadcast mechanism.

But the property of *being a global workspace* is *not* an all-or-none property. It's quite likely that many nonhuman animals *kind of* have a global workspace, but not *really*. So what should we say in those cases? Given that *phenomenal consciousness* is an all-or-none concept, we can't answer that question by saying that these nonhuman animals are *more or less* phenomenally conscious. For this reason (among others), Carruthers' response is that there's no answer. There's no fact of the matter.

Wouldn't it be possible to identify core global workspace features that can be determinately present or absent? I share Carruthers's skepticism on this question. It's not clear what global workspace features should be considered essential for phenomenal consciousness. It's not clear either how we could test for this in nonhuman animals if we don't already know whether they are phenomenally conscious or not (Michel 2019). And finally, it's not clear that the systems constituting core global workspace mechanisms such as attention or working memory would themselves come in determinate forms.

Instead, an obvious way of avoiding Carruthers's conclusion is to reject the global workspace theory. Carruthers suggests this possibility, but argues that "it can't be a condition of adequacy for a theory of human consciousness that it should be able to accommodate our intuitions about animals—especially since those intuitions vary widely across people" (144). That sounds right. But this answer is missing the point.

While intuitions may vary on the distribution of consciousness, they (probably) do *not* vary on the idea that *there is a fact of the matter*. As Jonathan Birch (forthcoming) points out, Carruthers seems to adopt a double standard. The intuition that consciousness is not graded appears to be beyond doubt, while the equally strong intuition that consciousness has to be determinate in all cases is rejected. If I were in a trolley problem deciding which intuition I should save, I, for one, would be ready to sacrifice the former to save the latter.

Fortunately, I'm not in this trolley problem. After all, is global workspace theory so strongly supported that we should be ready to drop the intuition that there's a fact of the matter about nonhuman animal consciousness? I don't think so. Let me explain.

Carruthers successfully defends the theory against the claim that prefrontal cortex activity—observed when contrasting conscious versus unconscious perception—only reflects cognitive requirements associated with reports (see also Michel and Morales 2019). But he fails to mention a more important confound in nearly all experiments interpreted as supporting global workspace theory: performance capacity (Lau 2008; Morales, Odegaard, and Maniscalco, forthcoming).

We usually find out about the neural correlates of consciousness by comparing conditions in which stimuli are consciously perceived to conditions in which they're unconsciously perceived. As predicted by the global workspace

theory, when this is done, consciousness seems to correlate with reverberating activity in a fronto-parietal global neuronal workspace network (Dehaene 2014). But that's not surprising. To make stimuli invisible, experimenters usually dampen the signal elicited by the stimulus to the point that it is barely discriminable. The signal is not similarly dampened in the conscious condition. So whatever contrast one gets from comparing conscious and unconscious conditions could reflect this signal strength difference—and associated differences in perceptual and cognitive capacities—rather than anything related to consciousness per se. We don't really know whether the evidence would still support a global workspace interpretation if this confound were avoided (although see Lau and Passingham 2006; Persaud et al. 2011).

Adding to this relative lack of unambiguous evidence is the fact that global workspace theory threatens to collapse under the weight of Carruthers's own theoretical construction.

Consider a humble bee. Its brain is the size of a sesame seed. This doesn't leave much room for global broadcast of the kind found in humans. But it could still enjoy a small-scale broadcast once in a while. According to Carruthers, there's no fact of the matter as to whether it is phenomenally conscious or not.

But now consider how Carruthers's take could backfire. My own brain has a very large number of such small-scale broadcasts, for instance in the dorsal visual stream, which Carruthers regards as “deeply inaccessible to consciousness” (57), or in the upper midbrain. Is there a fact of the matter as to whether *these* small-scale broadcasts sustain their own phenomenal consciousness? If yes, why can't we apply the exact same reasoning to the bee's small-scale global broadcast? If not, then there's no fact of the matter as to whether *any* mental activity happens unconsciously in humans. And if there's no fact of the matter about that, support for the theory vanishes, because there's no way to get any evidence that processing *outside* of our main global workspace—which happens to be hooked up with our theory of mind system—is phenomenally *unconscious*.

Whether or not you ultimately agree with the conclusion, Carruthers makes it clear that something has to give: either the global workspace theory, or some of our dearest intuitions about the nature of consciousness. As always, Carruthers is insightful and builds his arguments on his seemingly inexhaustible knowledge of the empirical literature. You'll certainly appreciate the journey.

References

- Birch, Jonathan. Forthcoming. “Global Workspace Theory and Animal Consciousness.” *Philosophical Topics*.
- Carruthers, Peter. 2006. *The Architecture of the Mind*. Oxford: Oxford University Press.
- Carruthers, Peter. 2011. *The Opacity of Mind*. Oxford: Oxford University Press.
- Carruthers, Peter. 2015. *The Centered Mind*. Oxford: Oxford University Press.

- Dehaene, Stanislas. 2014. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. New York: Penguin Books.
- Lau, Hakwan C. 2008. "Are We Studying Consciousness Yet?" In *Frontiers of Consciousness: Chichele Lectures*, edited by Lawrence Weiskrantz and Martin Davies, 245–58. Oxford: Oxford University Press.
- Lau, Hakwan C., and Richard E. Passingham. 2006. "Relative Blindsight in Normal Observers and the Neural Correlate of Visual Consciousness." *Proceedings of the National Academy of Sciences of the United States of America* 103, no. 49: 18763–68.
- Michel, Matthias. 2019. "Fish and Microchips: On Fish Pain and Multiple Realization." *Philosophical Studies* 176, no. 9: 2411–28.
- Michel, Matthias, and Jorge Morales. 2019. "Minority Reports: Consciousness and the Prefrontal Cortex." *Mind and Language* 35, no. 4: 493–513.
- Morales, Jorge, Brian Odegaard, and Brian Maniscalco. Forthcoming. "The Neural Substrates of Conscious Perception without Performance Confounds." In *Neuroscience and Philosophy*, edited by Felipe de Brigard and Walter Sinnott-Armstrong. Cambridge, MA: MIT Press.
- Persaud, Navindra, Matthew Davidson, Brian Maniscalco, Dean Mobbs, Richard E. Passingham, Alan Cowey, and Hakwan Lau. 2011. "Awareness-Related Activity in Prefrontal and Parietal Cortices in Blindsight Reflects More Than Superior Visual Performance." *NeuroImage* 58, no. 2: 605–11.

Matthias Michel

New York University

Philosophical Review, Vol. 130, No. 4, 2021

DOI 10.1215/00318108-9264069

J. P. Studd, *Everything, More or Less: A Defence of Generality Relativism*. Oxford: Oxford University Press, 2019. 304 pp.

Everything, More or Less by J. P. Studd is a joy to read. Its topic is the question of absolute generality: Is it possible to quantify over absolutely everything? This is a notoriously murky question. Much of the debate about absolute generality is about what the debate about absolute generality is about, and this book is no exception. Besides proposing a way of making sense of the question of absolute generality, it argues for *relativism*, the view that it is not possible to quantify over absolutely everything, and so against *absolutism*, the opposing view. What is exceptional about this book is how clearly and carefully it is written. It proceeds in an unhurried manner, without ever becoming boring or repetitious, and displays fastidious attention to detail throughout. The development and defense

Thanks to James Studd for helpful comments on a draft of this review.