Kinds of Consciousness
ECOM's 5th Graduate Conference

Expression, Communication, and Origins of Meaning Research Group

February 10, 2024

Registration Link: `https://ecomresearchgroup.com/kinds-of-consciousness/`

Conference Schedule

09:00 - 09:15 AM	Welcome and Introduction
09:15 - 10:25 AM	**Jonathan Birch: "The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI"**
10:35 - 11:05 AM	Christian de Weerd: "Sentience Gaps and Non-human Sentience Research"
11:10 - 11:40 AM	Jamie Freestone: "A three-part model of conscious awareness based on altered states"
11:45 - 12:15 PM	David R. Rowley: "Fading Qualia Arguments for Consciousness Supervening on Mathematical Structures"
12:15 - 01:00 PM	Lunch Break
01:00 - 01:30 PM	Alicia Peters: "Must an individual recognise their mental states as 'mine' in order to be phenomenally conscious? "
01:35 - 02:05 PM	Tiina Carita Rosenqvist: "Perceptual Competences and the Primary/Secondary Quality Distinction"
02:10 - 02:40 PM	Azenet Liora Lopez: "Split Brains and the Attentional Basis of Conscious Unity"
02:50 - 04:00 PM	**Elizabeth Schechter: "Self-Consciousness after Split-Brain Surgery"**

Abstracts

Jonathan Birch (LSE): "The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI"

Can octopuses feel pain and pleasure? What about crabs, shrimps, insects or spiders? How do we tell whether a person unresponsive after severe brain injury might be suffering? When does a fetus in the womb start to have conscious experiences? Could there even be rudimentary feelings in miniature models of the human brain, grown from human stem cells? And what about AI? These are questions about the "edge of sentience", and they are subject to enormous, disorienting uncertainty. The stakes are immense, and neglecting the risks can have terrible costs. We need to err on the side of caution in these cases, yet it's often far from clear what 'erring on the side of caution' should mean in practice. When are we going too far? When are we not doing enough? My forthcoming book *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI* constructs a precautionary framework designed to help us reach ethically sound, evidence-based decisions despite our uncertainty. This talk will introduce some of the main ideas, zooming in on the case of invertebrates.

Christian de Weerd (University of Groningen): "Sentience Gaps and Non-human Sentience Research"

Sentience is typically taken to be important in animal welfare considerations (Browning & Veit, 2022) and policy recommendations (e.g., Birch et al., 2021), and recent debates on non-human sentience have extended to artificial systems (Dung, 2023; Saad & Bradley, 2022; Shevlin, 2021; Müller, 2021; Ziesche & Yampolskiy, 2018; Metzinger, 2021). Within the context of non- human sentience, a prevalent way to conceptualize sentience is in a narrow sense picking out specifically *valenced conscious experiences* (Browning & Birch, 2021). These refer to those kinds of phenomenally conscious experiences that either feel good or bad, for examples states such as pain, fear, joy, pleasure, or relief (Birch et. 2021). Sometimes this is referred to as *affective* sentience (Powell & Mikhalevich, 2021). Unsurprisingly, most work on non-human sentience focusses on understanding, or developing indicators for, affective sentience (e.g., Crump et al., 2022; Sneddon et al., 2014; Elwood & Adams, 2015; Feinberg & Mallat, 2017).

However, appropriate sentience indicators should be sensitive to both consciousness *and* valency. One worry is that sentience indicators (see Browning & Birch, 2021), or frameworks incorporating several indicators (e.g., Sneddon et al., 2014; Crump et al., 2022), envisioned to target sentience directly mainly track affective states, leaving it unclear whether these states are actually conscious (Dawkins, 2012; 2017; 2021). The thought, then, is that such indicators might accurately track (the resulting behavior of) valenced states, yet such valenced states are simply processed unconsciously. The nomological possibility of permanently unconscious valenced states falls out of contemporary thinking and theorizing about valency. Many contemporary theories of valence conceptualize valency in terms of representational states representing non-mental properties like harmfulness (Cutter & Tye, 2014), value (Carruthers, 2017), (higher-order) imperatives (Klein, 2015; Martinez & Barlassina, forthcoming; Barlassina & Hayward, 2019) or desires to get rid of a sensation (Heathwood, 2007). It is taken to be a desideratum of valency accounts that they allow for unconscious valenced states (Carruthers, 2019). Similar points apply to functional accounts of valence (e.g., Damasio & Carvalho, 2013; Ginsburg & Jablonka, 2019).

Pragmatically, such worries are avoided by appealing to so-called precautionary principles (Birch, 2017). However, precautionary principles don't clarify the question of which systems are sen-

tient. Therefore, it's been suggested that insights on non-human affective states should be supplemented with insights on non-human consciousness (Browning & Birch, 2021; Shevlin, manuscript; Browning & Veit, 2023), a field that has seen a surge of new developments (e.g., Birch, 2022; Dung & Newen, 2023; Butlin et al., 2023; de Weerd, 2023; Shevlin, 2021). Non-human consciousness research has mainly focused on perceptual consciousness (e.g., Newen & Dung, 2023; Birch et al., 2021). These considerations lead to a natural, and often (implicitly) endorsed, view that accepts a (I will argue naïve) co-dependence between insights on non-human valency and consciousness. Specifically, to infer the presence of sentience is to simply combine the most plausible consciousness indicators and valency indicators (see e.g., Browning & Birch, 2022; Shevlin, manuscript; Browning & Veit, 2023). Emblematic of this way of thinking is the suggestion that non-human consciousness and non-human valency research should work in *parallel* (Shevlin, manuscript). The main criticism of approach is typically that consciousness is too ill-understood to make any meaningful predictions about (e.g., Dawkins, 2021).

However, even with a solid understanding of non-human consciousness, a problem that runs much deeper, caused by sentience gaps, remains. Sentience gaps occur in cases where the presence of phenomenal consciousness and valency in a system doesn't entail that the system is sentient. In such cases, one might correctly infer (using consciousness indicators) that a system is conscious (i.e., having visual conscious experiences), and correctly infer (using valency indicators) that a system has valenced states, yet wrongly conclude that the system is therefore sentient. In this case, a system might only have non-valenced perceptual conscious experiences, something taken to be possible (e.g., Carruthers, 2018; Barlassina & Hayward, 2019; Butlin et al. 2023; Godfrey-Smith, 2019), and only have non-conscious valenced states. Thus, even if conscious and valenced states are accurately tracked by consciousness and valency indicators respectively, this doesn't *necessarily* entail that a system is sentient. Sentience gaps are a residue, or artifact, of contemporary thinking and theorizing about consciousness and valency.

Sentience gaps are problematic because they challenge the immediate move of combining insights on non-human valency and non-human consciousness to the inference of sentience. Therefore, sentience gaps need to be considered when thinking about non-human sentience. How to deal with sentience gaps? I discuss at least two ways forward. Firstly, one might continue to accept the conceptual dissociation between consciousness and valency and simply put further constraints on their co-dependence. However, I argue that such responses are susceptible to more pernicious variants of sentience gaps caused by unclarity about how valency and consciousness *combine* into sentient states. Secondly, sentience gaps might not *in themselves* be problematic, but rather a symptom of inappropriate conceptualizing about consciousness and valency from the outset. If so, approaches that conceptually intertwine consciousness and valency should be taken more seriously (e.g., Damasio & Damasio, 2023; Veit, 2022).

References

- Barlassina, L., & Hayward, M. K. (2019). More of me! Less of me!: Reflexive Imperativism about Affective Phenomenal Character. *Mind, 128*(512), 1013-1044. `https://doi.org/10.1093/mind/fzz035`
- Berridge, K. C., & Kringelbach, M. L. (2015). Pleasure systems in the brain. *Neuron, 86*(3), 646-664. `https://doi.org/10.1016/j.neuron.2015.02.018`
- Birch, J. (2017). Animal sentience and the precautionary principle. *Animal Sentience, 16*(1). `https://doi.org/10.51291/2377-7478.1200`

- Birch, J. (2022). The search for invertebrate consciousness. *Noûs, 56*(1), 133–153. https://doi.org/10.1111/nous.12351
- Birch, J., Burn, C., Schnell, A., Browning, H., & Crump, A. (2021). Review of the evidence of sentience in cephalopod molluscs and decapod crustaceans. LSE Consulting. Available at: https://www.lse.ac.uk/News/News-Assets/PDFs/2021/Sentience-in-Cephalopod-Molluscs-and-Decapod-Crustaceans-Final-Report-November-2021.pdf
- Block, N. (1995). On a confusion about a function of consciousness. *Brain and Behavioral Sciences, 18*(2), 227-247. https://doi.org/10.1017/S0140525X00038188
- Browning, H., & Veit, W. (2022). The sentience shift in Animal Research. *The New Bioethics, 28*(4), 299-314. https://doi.org/10.1080/20502877.2022.2077681
- Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S. M., Frith, C., Ji, X., Kanai, R., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M. A. K., Schwitzgebel, E., Simon, J., & VanRullen, R. (2023). *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness.* arXiv. http://arxiv.org/abs/2308.08708
- Carruthers, P. (2018). Valence and Value. Philosophy and Phenomenological Research, 97(3), 658-680. https://doi.org/10.1111/phpr.12395
- Carruthers, P. (2020). Stop Caring about Consciousness. *Philosophical Topics, 48*(1), 1-20.
- Crump, A., Browning, H., Schnell, A., Burn, C., & Birch, J. (2022). Sentience in decapod Crustaceans: A general framework and review of the evidence. *Animal Sentience, 32*(1). https://doi.org/10.51291/2377-7478.1691
- Cutter, B., & Tye, M. (2011). Tracking Representationalism and the Painfulness of Pain. *Philosophical Issues, 21*(1), 90-109. https://doi.org/10.1111/j.1533-6077.2011.00199.x
- Damasio, A., & Damasio, H. (2023). Feelings Are the Source of Consciousness. *Neural Computation, 35*(3), 277-286. https://doi.org/10.1162/neco_a_01521
- Damasio, A., & Carvalho, G. B. (2013). The nature of feelings: evolutionary and neurobiological origins. *Nature Reviews Neuroscience, 14*, 143-152. https://doi.org/10.1038/nrn3403
- Dawkins, M. S. (2012). *Why animals matter: Animal consciousness, animal welfare, and human well-being.* Oxford University Press.
- Dawkins, M. S. (2017). Animal welfare with and without consciousness. *Journal of Zoology, 301*, 1-10. https://doi.org/10.1111/jzo.12434
- Dawkins, M. S. (2021). *The science of animal welfare: Understanding what animals want.* Oxford University Press.
- De Weerd, C. R. (2023). *A Credence-based Theory-heavy Approach to Non-human Consciousness.* PhilPapers. https://philarchive.org/rec/DEWACT
- Dung, L. (2023). How to deal with risks of AI suffering. *Inquiry.* https://doi-org.proxy-ub.rug.nl/10.1080/0020174X.2023.2238287

- Dung, L., & Newen, A. (2023). Profiles of animal consciousness: A species-sensitive, two-tier account to quality and distribution. *Cognition*, 235, 105409.
- Elwood, R. W., Adams, L. (2015). Electric shock causes physiological stress responses in shore crabs, consistent with prediction of pain. *Biology Letters, 11,* 20150800. https://dx.doi.org/10.1098/rsbl.2015.0800
- Feinberg, T. E., & Mallatt, J. M. (2017). *The Ancient Origins of Consciousness.* MIT Press.
- Ginsburg, S., & Jablonka, E. (2019). *The Evolution of the Sensitive Soul: Learning and the Origins of Consciousness.* MIT Press.
- Godfrey-Smith, P. (2020) *Metazoa: Animal Life and the Birth of the Mind*, New York: Farrar, Strauß, and Giroux.
- Gottlieb, J. (2022). *Capacity for Welfare and the Unity of Consciousness. Rethink Priorities.* Available at: https://docs.google.com/document/d/1OPkIgbvnylU-GF9M0ZxGxrVhlnySxLoq/edit#
- Heathwood, C. (2007). The Reduction of Sensory Pleasure to Desire. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 133*(1), 23-44.
- Klein, C. (2015). *What the body commands: the imperative theory of pain.* MIT Press.
- Lacewing, M. (2007). Do Unconscious Emotions Involve Unconscious Feelings. *Philosophical Psychology, 20*(1), 81-104. https://doi.org/10/10.80/09515080601023402
- LeDoux, J. E., & Brown, R. (2017). A higher-order theory of emotional consciousness. *Proceedings of the National Academy of Sciences of the United States of America, 114*(10), E2016-E2025. https://doi.org/10.1073/pnas.1619316114
- Martinez, M., & Barlassina, L. (forthcoming). The informational profile of valence: The metasemantic argument for imperativism. *British Journal for the Philosophy of Science.*
- Metzinger, T. (2021). Artificial Suffering: An Argument for a Global Moratorium on Synthetic Phenomenology. *Journal of Artificial Intelligence and Consciousness, 8*(1), 43-66. https://doi.org/10.1142/S270507852150003X
- Müller, V. C. (2021). Is it time for robot rights? Moral status in artificial entities. *Ethics and Information Technology, 23*, 579-587. https://doi.org/10.1007/s10676-021-09596-w
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review, 83*, 435-450. https://doi.org/10.2307/2183914
- Paul, E. S., Sher, S., Tamietto, M., Winkielman, P., & Mendl, M. (2020). Towards a comparative science of emotion: Affect and consciousness in humans and animals. *Neuroscience and Biobehavioral Reviews, 108,* 749-770. https://doi.org/10.1016/j.neubiorev.2019.11.014
- Power, R., & Mikhalevich, I. (2020). Affective sentience and moral protection. *Animal Sentience, 29*(35). https://doi.org/10.51291/2377-7478.1668
- Saad, B., & Bradley, A. (2022). Digital suffering: why it's a problem and how to prevent it. *Inquiry.* https://doi.org/10/1080/0020174X.2022.2067780

- Shevlin, H. (2021a). How Could We Know When a Robot was a Moral Patient? *Cambridge Quarterly of Healthcare Ethics, 30*(3), 459-471. https://doi.org/10/1017/S0963180120001012
- Shevlin, H. (2021b). Non-human consciousness and the specificity problem: a modest theoretical proposal. *Mind & Language, 36*(2), 297-314. https://doi.org/10.1111/mila.12338
- Shevlin, H. (2023). Negative valence, felt unpleasantness, and animal welfare. Available at: https://henryshevlin.com/wp-content/uploads/2021/11/Felt-unpleasantness.pdf
- Sneddon, L. U., Elwood, R.W., Adamo, S.A., & Leach, M.C. (2014). Defining and assessing Animal pain. *Animal Behaviour, 97*, 201-212. https://dx.doi.org/10.1016/j.anbehav.2014.09.007
- Veit, W. (2022). The Origins of Consciousness or the War of the Five Dimensions. *Biological Theory, 17,* 276-291. https://doi.org/10.1007/s13752-022-00408-y
- Velasco, P. F., & Loev, S. (2021). Affective experience in the predictive mind: a review and new integrative account. *Synthese, 198,* 10847-10882. https://doi.org/10.1007/s11229-020-02755-4
- Ziesche, S., & Yampolskiy, R. (2018). Towards AI Welfare Science and Policies. *Big Data and Cognitive Computing, 3*(1), 2. https://doi.org/10.3390/bdcc3010002

Jamie Freestone (Australian National University): "A three-part model of conscious awareness based on altered states"

Consciosuness is often divided into two separate and dissociable capacities. For example, phenomenal (P) consciousness versus access (A) consciousness (Block 1995); or sensation versus perception (Humphrey 2022; Reid 1785). A condition such as blindsight would seem to be a natural experiment that confirms the dissociability of these two capacities, at least for the visual modality (Block 1995; Humphrey 2022).

One recent motivation for a two-part model is that solving the hard problem of consciousness can be focused down to one component. A-consciousness might be conceivable in an autonomous system that nonetheless lacks any sense of "what it is like to be" that system: a system with generalised blindsight. And so P-consciousness — which involves the experiential, qualitative features of consciousness — would be the crucial thing to explain in order to solve the hard problem, rather than A-consciousness.

I endorse the rationale behind this framing as well as the focus on altered states of consciousness like blindsight. But I suggest that a three-part division is preferable. Specifically, A-consciousness needs to be replaced and P-consciousness can be su-divided between sensation (qualia in visual and other senses) and what I call the experiential component of conscious awareness. In this framing, full waking consciousness is the combination of three capacities:

1. concepts: being able to think in concepts, language, explicit categories;
2. sensation: the vivid, qualia-laden information from the senses; and
3. experience: the sense of continuity or ongoing witnessing of experience.

This can be illustrated by altered states of consciousness in which one or two components are present but not all three. Many altered states of consciousness are not well studied and often absent from work in philosophy of mind. I look at the nascent literature on four altered states that together demonstrate the three-part model I propose: non-rapid-eye-movement (NREM) sleepwalking, ego-dissolution with psychedelic drugs, lucid dreamless sleep, and what I'm calling *sleep thinking*.

NREM sleepwalkers lack *experience*. The literature is divided over exactly how conscious they are. The assumption used to be that NREM sleepwalkers are utterly unconscious, yet able to perform complex motor tasks and rote exercises (Zadra & Pilon 2012). Some recent work suggests that they might experience momentary inchoate flashes of sensation, but without ongoing awareness (Rocha & Arnulf 2020).

Lucid dreamless sleep is an ancient practice only recently named. A subject may, through much training in meditation, achieve a kind of pure awareness free of any conceptual or sensual objects. Some people also report reaching this via sleep (Alcara-Sanchez 2022). They become lucid, in the sense that they are aware of their own experience and remember it, but it has no "content".

During **ego-dissolution** or ego-death, the subject will maintain consciousness inasmuch as they remember the episode and during the episode have a sense of experiencing something, often very vivid sensations and imagery. Typically, however, they cannot name or classify the objects of that experience during the trip itself and their sense of self or ego dissolves (Lebedev *et al.* 2015). Their linguistic and conceptual abilities temporarily shutdown. The experience is like a temporary version of what some stroke victims experience when parts of the brain's left-hemisphere are damaged (Taylor 2009).

Sleep thinking is perhaps a fairly common experience but is only recently described in any academic literature and then only as a side note to a discussion of the also relatively under-theorised lucid dreamless sleep (Alcarez-Sanchez 2023; Metzinger 2020). In sleep thinking, the subject becomes aware and can think in a non-surreal, deliberate manner. Semantic memories are available and so are words and concepts. But the sleep thinker is in a void, bereft of imagery or sensory input of any kind: a kind of globalised aphantasia. Sleep thinking is lucid dreamless sleep with voluntary conceptual thought.

Opinions will differ over which of these states deserves to be called "conscious". Personally, I think lucid dreamless sleep, sleep thinking, and ego dissolution are; NREM sleepwalking isn't. Regardless of interpretation, one can acknowledge that these states decouple from one another the traditional dyad of A-consciousness and P-consciosuness and indeed the three components I have suggested. Sleep thinkers have concepts, no sensations. Drug trippers have sensation but no concepts. And these two have something in common which NREM sleepwalkers lack, even if they have sensation and/or concepts — this crucial component is *experience* which can also be isolated from sensation and concepts, as in the MPE of lucid dreamless sleep.

What I'm calling *experience* may well align with pre-existing ideas in consciousness research. Perhaps it is identical to the global workspace (Baars 2005), or the "attention schema" (Graziano 2020), or a form of working memory. In any case, it appears to be a necessary condition for what we'd normally call conscious awareness. I suspect it is also sufficient: my theory predicts that most people who have an MPE will consider it to be a conscious experience even though it lacks so-called qualia and conceptual thought. This shifts the focus of the hard problem to explaining experience rather than sensation or qualia — which may be a separate and equally hard problem, but is not *the*

hard problem that is the target of illusionism and other recent theories of consciousness (Frankish 2019).

The theory also predicts that at least some NREM sleepwalkers should be able to perform tasks that require sensation or concepts. The question then becomes: which tasks can NREM sleepwalkers not perform? Equivalently, what are the tasks which require the experience component of consciosuness? My speculation is that these would be certain social cognition tasks which unfold over overlapping timescales, and hence require an ongoing, narrative-like stream of consciousness, beyond flashes of sensation. Examples might include caretaking of infants, cooperating on a project, and shared intentionality based tasks (Tomasello & Moll 2010).

References

- Alcaraz-Sánchez, A. (2022). Is Lucid Dreamless Sleep Really Lucid?. *Review of Philosophy and Psychology*, 1-27.
- Alcaraz-Sánchez, A. (2023). Awareness in the void: a micro-phenomenological exploration of conscious dreamless sleep. *Phenomenology and the Cognitive Sciences, 22*(4), 867-905.
- Baars, Bernard J. "Global workspace theory of consciousness: toward a cognitive neuroscience of human experience." *Progress in brain research 150* (2005): 45-53.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and brain sciences, 18*(2), 227-247.
- De Haan, E. H., Corballis, P. M., Hillyard, S. A., Marzi, C. A., Seth, A., Lamme, V. A., ... & Pinto, Y. (2020). Split-brain: what we know now and why this is important for understanding consciousness. *Neuropsychology review, 30*, 224-233.
- Frankish, K. (2019). The meta-problem is the problem of consciousness. *Journal of Consciousness Studies, 26*(9-10), 83-94.
- Freestone, NREM (forthcoming)
- Graziano, M. S. (2020). Consciousness and the attention schema: Why it has to be right. *Cognitive Neuropsychology, 37*(3-4), 224-233.
- Humphrey, N. (2022). *Sentience: The Invention of Consciousness.* Oxford University Press.
- Lebedev, A. V., Lövdén, M., Rosenthal, G., Feilding, A., Nutt, D. J., & Carhart-Harris, R. L. (2015). Finding the self by losing the self: Neural correlates of ego-dissolution under psilocybin. *Human brain mapping, 36*(8), 3137-3153.
- Metzinger, T. (2020). Minimal phenomenal experience: Meditation, tonic alertness, and the phenomenology of "pure" consciousness. *Philosophy and the Mind Sciences, 1*(I), 1-44.
- Nour, M. M., Evans, L., Nutt, D., & Carhart-Harris, R. L. (2016). Ego-dissolution and psychedelics: validation of the ego-dissolution inventory (EDI). *Frontiers in human neuroscience*, 269.
- Reid, T. (1850). *Essays on the intellectual powers of man.* J. Bartlett.
- Rocha, A. L., and I. Arnulf. "NREM parasomnia as a dream enacting behavior." *Sleep medicine, vol. 75*, 2020, pp. 103-105.

- Taylor, J. B. (2009). *My stroke of insight.* Hachette UK.
- Tomasello, M., & Moll, H. (2010). The gap is social: Human shared intentionality and culture. *Mind the gap: Tracing the origins of human universals*, 331-349.
- Zadra, A. and Pilon, M. (2012). *Parasomnias II: Sleep terrors and somnambulism.* The Oxford Handbook of Sleep and Sleep Disorders.

David R. Rowley (University of Leeds): "Fading Qualia Arguments for Consciousness Supervening on Mathematical Structures"

Consciousness is simultaneously familiar yet elusive. As it appears qualitatively different from other parts of existence, any complete account of reality will need to account for consciousness. An understanding of what consciousness is could also give indications of how it is connected with other parts of reality, in particular physical objects.

Chalmers' fading qualia argument (Chalmers, 1996, pp. 253–263) considers incrementally changing a consciousness-supporting brain into a physically different, but functionally equivalent, object. As it is implausible that any particular single small change would affect the existence of the consciousness, Chalmers suggests that the functionally equivalent object would also have consciousness, and hence that conciousness does not supervene on physical objects, but it could supervene on functions. This paper applies similar arguments on structural isomorphs to argue that consciousness is a property of mathematical structures.

Assuming that consciousness supervenes on the brain in some way, this paper firstly considers whether it supervenes on physical objects, structures or functions. Functions are considered to differ from structures in that functions involve changes to states, via dependency relations, while structures could be static. Chalmers' fading qualia argument is applied to the transformation of a brain into a structurally isomorphic, but functionally and physically different, object. One way to do this would be to change, one by one, the dependency relations of individual components in the brain. For example, dependency relations could be changed from the form 'neuron x fires if and when neuron y fires' to a form 'neuron x fires at times t_1, t_2, etc.' in such a way that results in a structurally similar object that behaves similarly. In line with the fading qualia argument, this can be done iteratively. Similarly, it would be implausible that any individual change would affect the existence of consciousness, suggesting that the final object would also be conscious. However, with very different dependency relations it would be functionally different. This suggests that consciousness cannot supervene on functions or physical objects, but could supervene on structures. This could be on physical instantiations of the structure, or the mathematical object itself.

Another argument for the experiential equivalence over structural isomorphs considers the speed that electrons and other relevant neurotransmitters move through the brain. We would not expect any incremental slowing of the movement of all components to change the experiential properties at that time (although it would be longer before these properties change). Continuing to slow down the movements without experiential change, suggests that we could eventually bring them to a stop without destroying the experience. We would then have an object that was structurally isomorphic to the brain but functionally inert. If such an object did indeed have experience then functionality is irrelevant and only the structure is important to consciousness.

The paper next considers gradually dividing a conscious object into two structurally isomorphic distinct physical objects. This can be done incrementally by replacing a single neuron with two

artificial neurons connected to the same point in the network. By continuing to replace single neurons with two artificial neurons and connecting them, the original object can gradually be transformed into two structurally isomorphic objects. The fading qualia argument can be applied again, suggesting that the number of consciousnesses would not vary over the transformation. This argues that if distinct objects are structurally isomorphic then the consciousnesses associated with the two systems would be numerically identical. There being two physical instantiations of the structure, but just the one mathematical object, suggests that consciousness supervenes on mathematical structures, rather than on physical instantiations.

A further experiment involves delaying one artificial neuron. If we replaced a single natural neuron with two artificial neurons we could introduce a delay to one of them. If artificial neurons respond more quickly than natural neurons, then one could briefly wait before reacting to stimuli. If it was set up that by the time the signal returns to the brain the two sources were synchronized then this should make no difference to the consciousness. We could then continue from here to two artificial brains, again saying (by the fading qualia argument) that they still involve exactly one consciousness. But now we have the two structures instantiated at different locations and times, but still relating to just the one consciousness. This suggests that consciousness is not related to something spatio-temporal (such as matter) but something outside of space and time (such as mathematical objects).

These arguments suggest that transforming a single physical instantiation of a structure into two instantiations does not produce another consciousness. Therefore, the consciousnesses across all structurally isomorphic physical objects are numerically identical. There being multiple physical objects, but only one abstract mathematical object, associated with the structure, indicates that consciousness is associated with the mathematical object, rather than the physical instantiations. Physical objects would then be causally connected to the consciousness via the mathematical structure they instantiate.

This model requires that mathematical objects exist, which under Full-Blooded Platonism means that all possible mathematical objects exist (e.g. Balaguer 1998, p5). As all consciousnesses consistent with each mathematical object exist in this model, then all possible consciousnesses exist, whether or not any physical object instantiates the corresponding structure. The model therefore gives an account of what consciousness is, and asserts that each possible conscious state exists exactly once, independently of physical objects.

References

- Balaguer, M. (1998). *Platonism and anti-Platonism in mathematics.* New York: Oxford University Press.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory.* Oxford: Oxford University Press.

Alicia Peters (University of Nottingham): "Must an individual recognise their mental states as 'mine' in order to be phenomenally conscious?"

In this paper, I aim to question whether the subjective and phenomenally conscious nature of consciousness is to be constituted by a sense of mine-ness that allows our mental states to be recognised as my own, or rather *mine*. With this in mind, I shall discuss the subjective nature of our mental states as they relate to the disordered phenomenon of "Thought Insertion".

Thought Insertion has become widely recognised as a diagnostic symptom of Schizophrenia, defined within the DSM-V (American Psychiatric Association, 2013: 208) as a delusion believed by schizophrenic individuals that their thoughts are not their own, but rather are inserted into their mind. I wish to integrate Thought Insertion into the discussion of the subjective nature of consciousness by posing the following question: Does the notion of Thought Insertion preclude an individual from possessing phenomenally conscious mental states?

So, what exactly does it mean to possess phenomenally conscious mental states? Well, Phenomenal Consciousness refers to the feeling of what it is like to be you (Nagel, 1974: 436), which is sometimes exemplified by the capacity to recognise an individual's mental states as mine: as possessing a sense of mine-ness.

This paper will investigate how Thought Insertion undermines the universality of all our mental states possessing a sense of mine-ness, which demonstrates how Phenomenal Consciousness should not always be interpreted in terms of a sense of mine-ness. My argument goes as follows: the sense of mine-ness is not a universal aspect of mental states, and so inserted thoughts remain phenomenally conscious despite the individual's lack of recognition of mine-ness. Rather, to appropriately capture the subjective and phenomenal nature of consciousness, we must realise that not all mental states are to be recognised by an individual as their own, *as mine*, to be phenomenally conscious.

A crucial element of my argument hinges on the erroneous assumption that an individual's possession of a sense of mine-ness constitutes their mental states being phenomenally conscious. Suppose a given individual: Susan, will have the thought, 'there is water in this glass.' As Ned Block, (1995: 228) puts it, what makes a state phenomenally conscious is that there is something it is like for an individual to be in it. Thus, this thought is phenomenally conscious so long as there is some unique, *Susan-esque*, way it feels like for Susan to experience it. Philosophers such as Guillot (2016) assume that this phenomenal character is to be characterised as a sense of mine-ness. So, mental states are phenomenally conscious so long as they can be identified as having a sense of mine-ness.

What do I mean by mine-ness (Guillot, 2016: 31)? Well, mine-ness refers to a specific construal of the subjectivity of experience which states that a subject's experiential awareness is of herself as related to the experience she is having. By this I mean, that a subject will be able to recognise her mental state as her own, so long as she is aware of her sense of self as related to the experience of this mental state. I take 'sense of self' to include who we are, our ambitions, and our psychological narrative. So, in this way, the assumption is that an individual can recognise her thought as her own, and this constitutes the fact that there is something it feels like for her to have this thought. In essence, her mental state is phenomenally conscious and her own, in virtue of the fact that she can recognise her sense of self as intimately related to her experience.

The idea that our mental states are phenomenally conscious due to a sense of mine-ness is thought to be a universal characteristic of our mentality, just as phenomenal consciousness is. However, there are times where this proves not to be the case. Enter Thought Insertion. Thought Insertion is a symptom of schizophrenia in which the individual believes that someone else's thoughts have been inserted into their mind, causing them to believe their thoughts are not their own. This is a renowned example of a case where an individual lacks a sense of mine-ness over their thoughts.

Consider this example: 'she said that sometimes it seemed to be her own thought ... but the feeling isn't the same, the feeling is that it is somebody else's (Allison-Bolger, 1999: #89).' Suppose this individual has the thought that 'I want to kill my sister.' This thought does not only feel inserted into her mind, but it also does not feel like her own, thus renouncing her sense of mine-ness. This individual will not feel her sense of self as related to the experience of this inserted thought she is having. Cases like this express how it is not a universal principle that all mental states must possess a sense of mine-ness for the individual experiencing them.

However, my argument is that we should forgo the assumption that our mental states are phenomenally conscious in virtue of a sense of mine-ness, thus differentiating between Phenomenal Consciousness and a sense of mine-ness. I believe that the schizophrenic individual nonetheless experiences their thoughts (and various other mental states) as phenomenally conscious, regardless of their lack of a sense of mine-ness. There is still something it is like for the individual to experience their thoughts, as not their own! So, we must disrupt the connection between mine-ness and phenomenal consciousness by realising that all individuals experience their mental states in a phenomenally conscious manner even if they cannot recognise their mental states as mine, or rather their own. Thus, Phenomenal Consciousness is a pervasive and ingrained feature of our mentality that cannot be disrupted or impaired, although a sense of mine-ness can be.

This paper is heavily motivated by my desire to expand the study of consciousness beyond the purely non-pathological setting and into the context of mental health. In investigating the nature of our mental states and consciousness through the lens of mental disorders, we can ameliorate our understanding of the subjective nature of mentality. My argument establishes how subjects of Thought Insertion experience their mental states as phenomenally conscious despite a lack of possession of a sense of mine-ness. I believe that my paper will demonstrate how what it is like to experience a given mental state need not be confined to the capacity to recognise a mental state as one's own, *as mine*.

References

- American psychiatric association. 2013. Schizophrenia spectrum and other psychotic disorders. In *Diagnostic and statistical manual of mental disorders: DSM-5.* (5th edition). Washington, D.C.: American Psychiatric Publishing.
- Nagel, T. 1974. What is it like to be a bat? In *Philosophical Review. 83* (October): 435-450
- Block, N. 1995. On a confusion about a function of consciousness. In *Brain and Behavioral Sciences. 18*(2): 227-247
- Guillot, M. 2016. I Me Mine: On a confusion concerning the subjective character of experience. In *Review of Philosophy and Psychology.* (1): 1-31
- Allison-Bolger, V. Y. 1999. 'Collection of case histories. Unpublished typescript. #89 In C, Hoerl. 2001. On thought insertion. *Philosophy, Psychiatry, and Psychology. 8*(2/3): 189-200

Tiina Carita Rosenqvist (Dartmouth): "Perceptual Competences and the Primary/Secondary Quality Distinction"

My goal in this talk is to re-conceptualize the Early Modern distinction between primary and secondary qualities. The original distinction is metaphysical: primary qualities (such as shapes

and sizes) are intrinsic or "real" properties of external objects, whereas secondary qualities (such as colors and smells) are perceiver-dependent in some way (*e.g.*, Galileo 1623, Locke 1689-90). My re-conceptualization is epistemological. I propose that we conceptualize primary qualities as the kind of properties that we can competently perceive and/or have perceptual knowledge of and secondary qualities as the kind of properties that play an enhancement role in the competent perception and/or perceptual knowing of primary qualities. For simplicity's sake, I focus on human perceivers.

The notion of 'perceptual competence' is important for my view. Though there are different philosophical accounts of competences (*e.g.*, Greco 2007, Sosa 2015, Miracchi 2015), I will stay neutral between those accounts.[1] I find the notion of 'perceptual competence' useful for three reasons. First, it captures the apparent skillfulness of perception, *i.e.*, the idea that perceptual capacities are often reliable. Second, it captures the idea that perception is ultimately *for* action; what counts as a competence for a particular type of animal depends on the action goals of that animal. Third, it captures the idea that perceptual capacities often involve skill and can be developed or gradually refined. For example, shape perception competence develops in childhood. By the age of ten, most children no longer underestimate the sizes of objects in the distance (Granrud 2012).

My proposal is that primary qualities are the kind of properties that we can competently perceive and/or have perceptual knowledge of. For example, let us say that a person is a competent perceiver of distances when her distance perceptions/judgements reliably correspond to reality at a scale required for successful action, *e.g.*, when an object at a distance of three meters will reliably look to be further away than an object at a distance of two meters, enabling successful interaction (approaching, grasping, avoiding, etc.). If we can establish that human distance is reliable in this way, we can label distance a primary quality.

Now let us consider shape. We can say that a person possesses a shape perception competence when square objects reliably look (to be) square-shaped and circular objects circular to her, etc. Vision science tells us that that humans are adept at judging objective shapes, at least if there is sufficient 3D context information available (*e.g.,* Sereno et al. 2020, Morales et al. 2020) and especially if they are allowed to view the object from different perspectives (*e.g.*, Wang et al. 2018). In addition, representations of objective shapes "form automatically and obligatorily" (Baker & Kellman 2020). These points are corroborated by phenomenology. Our spatial experience appears to be geared toward reality. In other words, we care a great deal about the objective shapes of things, not (just) their perspectival appearances. That said, the way those appearances behave is likely important. An object will generally appear the same way from the same visual angle and this predictability makes competent shape perception possible. And since we reliably perceive / perceptually know the objective shapes of things in the world, shape is a primary quality.

Not all types of perceptual processing are reliable in this way. For example, color visual processing does not issue in perceptions that neatly correspond to reality. Variation in color perception is systematic and rife; color perceptions are influenced by lighting conditions[2], chromatic contexts, viewing angles, viewing distances, etc. The best explanation for such variation (as I argue in detail elsewhere) is that color perception is not itself a competence but that color visual processing is 'competence-embedded.' More specifically, I propose that the function of color vision is to enable and enhance the manifestation of a variety of different perceptual competences, such as scene segmentation, shadow disambiguation, and property identification, not to enable the perception of

[1] My use of the term 'perceptual competence' is liberal in the sense that it is meant to cover both purely perceptual (*i.e.*, non-cognitive) competences and perceptual-cognitive competences.

[2] Note that color constancy is approximate at best (*e.g.*, Shevell & Kingdom 2008; Foster 2011). This applies to both the constancy of phenomenal experience across lighting conditions and the constancy of perceptual judgements concerning the material properties of surfaces across lighting conditions.

color *per se*. Our perceptual experience also doesn't seem to reach beyond color appearances in the same way that it reaches beyond shape appearances. Nor do color appearances enable the kind of extrapolation that shape appearances enable; color is unpredictable. And so, because color perception is competence- embedded (and not a competence in its own right), color is a secondary quality.

On a more general level the view can be summarized as follows: some perceptual properties are what we perceive and others are how we perceive. With this framework in place, we can tell a coherent and explanatorily robust story about human perceptual processing, and make sense of a wide variety of perceptual phenomena. For example, I have proposed that humans aren't competent perceivers of color, but color helps us competently perceive shapes. This helps explain the existence of simultaneous contrast effects and color assimilation effects in color perception (as illustrated by many "textbook color illusions"). Interestingly, I suspect that even the primary qualities thus conceived will turn out to include at least some perceiver-dependent properties. I therefore reject the widely held view that sensory perception is in the business of presenting or representing the objective world as it is, and side with those philosophers and scientists who argue that the fundamental goal of perception is to be a guide to adaptive behavior and that perceptions do not always have to align with physical reality to successfully play this role (*e.g.*, Hatfield 2009, Purves et al. 2015, Hoffman et al. 2015).

References

- Baker N, Kellman PJ (2018). Abstract shape representation in human visual perception. *Journal of Experimental Psychology: General 147*(9): 1295–1308.
- Galileo G (1623). Il Saggiatore (The Assayer), Rome. Parts translated in Maurice A. Finocchiaro (Ed. and Trans.), *The Essential Galileo*. Indianapolis: Hackett, 2008, 179–189.
- Granrud CE (2012). Judging the size of a distant object: Strategy use by children and adults. In G Hatfield & S Allred (Eds.), *Visual experience: Sensation, cognition, and constancy*. Oxford University Press, 13–34.
- Greco J (2007) The nature of ability and the purpose of knowledge. *Philosophical Issues 17*:57–69.
- Hatfield, G. (2009). On Perceptual Constancy. In G. Hatfield, *Perception and Cognition: Essays in the Philosophy of Psychology*. Oxford: Clarendon Press, 189-211.
- Hoffman DD, Singh M, Prakash C (2015). The Interface Theory of Perception. *Psychonomic Bulletin* & Review 22:1480–1506.
- Locke J (1689–90 [1975]). *An Essay concerning Human Understanding*, London. Reprinted with Peter H. Nidditch (Ed.). Oxford: Clarendon Press.
- Miracchi L (2015) Competence to Know. *Philosophical Studies 172*: 29–56.
- Morales J, Bax A, Firestone C (2020). Sustained representation of perspectival shape. *PNAS 117*(26): 14873–14882.
- Purves, D., Morgenstern, Y., & Wojtach, W.T. (2015). Perception and Reality: Why a Wholly Empirical Paradigm is Needed to Understand Vision. *Frontiers in Systems Neuroscience 9:* 156.

- Sereno ME, Robles KE, Bies AJ (2020). The Effects of Three-Dimensional Context on Shape Perception. *Psychological Science 31*(4): 381–396.
- Shevell SK & Kingdom FAA (2008). Color in complex scenes. *Annual Review of Psychology 59*: 143–166.
- Sosa E (2015) *Judgment and Agency*. Oxford: Oxford University Press.
- Wang XM, Lind M, Bingham GP (2018). Large Continuous Perspective Change With Non-coplanar Points Enables Accurate Slant Perception. *Journal of Experimental Psychology: Human Perception and Performance 44*(10): 1508–1522.

Azenet Liora Lopez (Ludwig Maximilian University of Munich): "Split Brains and the Attentional Basis of Conscious Unity"

For decades, split-brain subjects have been a valuable test case for many theories and claims about consciousness. Though they report feeling normal and behave normally in their day-to-day lives (Sperry 1982; Pinto, Neville et al. 2017), under experimental conditions each of their disconnected hemispheres seems to bear conscious contents not shared by the other hemisphere.[3] Most notably, split-brain subjects can accurately identify pairs of objects presented to opposite hemispheres but cannot make same/different judgments about these objects (Johnson 1984; Sergent 1986; Fendrich and Gazzaniga 1989; see also Pinto, Neville et al. 2017). Such findings have led psychologists, neuroscientists, and philosophers to believe that split-brain consciousness might be split into two independent streams (Sperry 1982; Volz and Gazzaniga 2017; Schechter 2018). Moreover, two-stream views are in accordance with current leading theories regarding large-scale integration of information across the brain as necessary for consciousness (such as the Global Neuronal Workspace Theory or the Integrated Information Theory; see Dehaene 2014, Tononi and Koch 2015).

To resist two-stream views, theorists have attributed split-brain subjects' cross-matching inability to disintegrations in processes other than consciousness itself, for instance, *perception* (Pinto, Neville et al. 2017; see also Sergent 1986) or conscious *contents* (Bayne 2008). However, existing one-stream models of split-brain consciousness do not yet account for one central kind of conscious unity, namely, *phenomenal unity*. According to an influential characterization (Bayne and Chalmers 2003), two contents of a subject S's experience, E1 and E2, are phenomenally unified if and only if E1 and E2 have a conjoint phenomenology, that is, if and only if there is something it is like for S to experience E1 and E2 together. Though existing one-stream models explain how right- and left-hemisphere conscious contents (call them 'RE' and 'LE') could inform the actions of a single conscious agent or even be the experiences of a single conscious subject, they do not specify how RE and LE could have a conjoint phenomenology.[4] One model, the *unified agency model* (Pinto, Neville et al. 2017; Pinto, de Haan et al. 2017), claims that RE and LE belong to two disintegrated streams of perception of a single conscious agent. Another model, the *switch model* (Bayne 2008, 2010), claims that RE and LE are sequentially activated and are almost never conscious at the

[3] Some argue that experimental conditions in fact disrupt an otherwise unified stream of consciousness. However, this view may require adopting the controversial externalist view that mental states are only fully individuated in terms of their environment. See recent discussion in Downey (2018).

[4] For discussion of how agent unity and subject unity could come apart from phenomenal unity, see Schechter and Bayne (2021), Dainton (2014).

same time. These claims are hard to reconcile with the possibility of RE and LE being experienced together; thus, on these models split-brain consciousness may still remain phenomenally disunified.

This paper offers an account of how split-brain consciousness could be phenomenally unified. My argument has two main premises:

1. Attentional capacities afford phenomenal unity: If S can exercise appropriate attentional capacities, S's consciousness is phenomenally unified.

2. Split-brain subjects can exercise appropriate attentional capacities.

Premise 1 rests on recent conceptual work by Watzl (2014, 2017) and Wiese (2022). According to Watzl, attention unifies the field of conscious experience by organising its elements into what is more central (i.e., in the focus of attention) or more peripheral (i.e., outside the attentional focus). This organisation creates a conjoint phenomenology between experiential elements because what it is for an element to be experienced as more central is for others to be experienced as more peripheral. In turn, Wiese argues that attention generates a phenomenally felt connection between any two elements of experience whenever we become aware of our ability to voluntarily shift our focus of attention from one element to the other. If these proposals are in the right track, then attention is indeed a plausible ground for the kind of conscious unity we are after.

Regarding premise 2, it is noteworthy that research on split-brain attention is somewhat mixed (Handy and Gazzaniga 2005; Pinto, de Haan et al 2017). The attentional system in humans is in fact constituted by a collection of distinct functions (Petersen and Posner 2012), and some of these do exhibit notorious breakdowns in the split brain. For example, in early processing stages, the right and left hemispheres seem to have independent foci of attention (Mangun et al. 1994). Shifts of attention across the midline take significantly longer than shifts within an hemifield, even though the distance is the same (Reuter-Lorenz and Fendrich1990). And split-brain subjects can also split attention to perform two simultaneous and conflicting tasks (Ellenberg & Sperry 1979, 1980). I argue that though these breakdowns do compromise phenomenal unity, insofar as they may generate separate centrality systems or impair attentional agency, other attentional capacities retain sufficient unity to support what I call a *layer* of phenomenal unity in split-brain consciousness. See Table 1.

Attentional function	**Breakdown?**
Exogenous capture	Severe
Top-down orienting	Mild (unified under appropriate conditions)
Executive control	Mild (unified under appropriate conditions)
Alertness	No

Table 1: Attention in the split brain

In the view I recommend, only *local* aspects of split-brain experiences are phenomenally disintegrated. These are constituted by the relatively detailed properties of individual objects experienced in both sides, such as their colours and shapes. However, this is no impairment for subjects exercising an amount of attentional agency over these experiential elements, nor for the latter to be connected in a single centrality system. I propose that this is achieved in virtue of the *global* aspects of these experiences, such as occupied locations and configurational properties. I call this proposal "the layered unity model".[5]

[5]Though I situate the view in the one-stream camp, it may depart from other one-stream models in that it allows

This model is a plausible way to make sense of the observation that the disconnected hemispheres share only crude or limited information, with low spatial and temporal resolution (Corballis 1995). It also aligns nicely with current approaches distinguishing two levels or aspects of phenomenal consciousness, based on cortical and subcortical structures (Newen & Montemayor 2023). Finally, if the layered unity model accurately captures the structure of split-brain consciousness, then theories entailing a full-blown two-stream view (such as the Integrated Information Theory or the Global Neuronal Workspace theory) might require some adjustments and/or refinements.

References

- Bayne, T. J. (2008). The unity of consciousness and the split-brain syndrome. *Journal of Philosophy, 105*(6), 277-300. https://doi.org/10.5840/jphil2008105638
- Bayne, T. J. (2010). *The Unity of Consciousness*. Oxford, GB: Oxford University Press UK.
- Bayne, T. J., & Chalmers, D. J. (2003). What is the unity of consciousness? Re-printed in Chalmers, D. J., *The Character of Consciousness*, 497–539. Oxford: Oxford University Press.
- Corballis, M. C. (1995). Visual integration in the split brain. *Neuropsychologia, 33*, 937–959. https://doi.org/10.1016/0028-3932(95)00032-X
- Dainton, B. (2014). Unity, synchrony, and subjects. In D. Bennett & C. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 255-286). Cambridge, MA: MIT Press.
- Dehaene, S. (2014). *Consciousness and the brain: Deciphering how the brain codes our thoughts*. Viking.
- Ellenberg, L., & Sperry, R. W. (1979). Capacity for holding sustained attention following commissurotomy. *Cortex; A Journal Devoted to the Study of the Nervous System and Behavior, 15*(3), 421–438. https://doi.org/10.1016/s0010-9452(79)80068-4
- Ellenberg, L., & Sperry, R. W. (1980). Lateralized division of attention in the commissurotomized and intact brain. *Neuropsychologia, 18*(4-5), 411–418. https://doi.org/10.1016/0028-3932(80)90144-x
- Fendrich, R., & Gazzaniga, M. S. (1989). Evidence of foveal splitting in a commissurotomy patient. *Neuropsychologia, 27*, 273-281.
- Handy, T. C., & Gazzaniga, M. S. (2005). Attention in split-brain patients. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of Attention*.
- Johnson, L. E. (1984). Bilateral visual cross-integration by human forebrain commissurotomy subjects. *Neuropsychologia, 22*(2), 167–175. https://doi.org/10.1016/0028-3932(84)90059-9
- Mangun, G. R., Luck, S. J., Plager, R., Loftus, W., Hillyard, S. A., Handy, T., Clark, V. P., & Gazzaniga, M. S. (1994). Monitoring the visual world: Hemispheric asymmetries and subcortical processes in attention. *Journal of Cognitive Neuroscience, 6*(3), 267–275. https://doi.org/10.1162/jocn.1994.6.3.267

split-brain consciousness to be disunified to some extent. In this respect, the view could be more aptly classified as a partial unity model (Schechter and Bayne 2021; Schechter 2014; Bayne 2010).

- Newen, A., & Montemayor, C. (2023). The ALARM Theory of Consciousness: A Two-Level Theory of Phenomenal Consciousness. *Journal of Consciousness Studies, 30*(3), 84-105.
- Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual Review of Neuroscience, 35*, 73–89. https://doi.org/10.1146/annurev-neuro-062111-150525
- Pinto, Y., Neville, D. A., Otten, M., Corballis, P. M., Lamme, V. A. F., de Haan, E. H. F., Foschi, N., & Fabri, M. (2017). Split brain: Divided perception but undivided consciousness. *Brain, 140*(5), 1231-1237.
- Pinto, Y., de Haan, E. H. F., & Lamme, V. A. F. (2017). The split-brain phenomenon revisited: A single conscious agent with split perception. *Trends in Cognitive Sciences, 21*(11), 835–851. https://doi.org/10.1016/j.tics.2017.09.003
- Reuter-Lorenz, P. A., & Fendrich, R. (1990). Orienting attention across the vertical meridian: Evidence from callosotomy patients. *Journal of Cognitive Neuroscience, 2*(3), 232–238. https://doi.org/10.1162/jocn.1990.2.3.232
- Schechter, E. (2014). Partial unity of consciousness: Evidence and implications. In D. J. Benett & C. Hill (Eds.), *The Unity of Consciousness and Sensory Integration* (pp. 347–373). MIT Press.
- Schechter, E. (2018). *Self-Consciousness and Split Brains: The Minds' I*. Oxford, UK: Oxford University Press.
- Schechter, E., & Bayne, T. J. (2021). Consciousness after split-brain surgery: The recent challenge to the classical picture. *Neuropsychologia, 160*, 107987. https://doi.org/10.1016/j.neuropsychologia.2021.107987
- Sergent, J. (1986). Subcortical coordination of hemisphere activity in commissurotomized patients. *Brain: A Journal of Neurology, 109*(Pt 2), 357–369. https://doi.org/10.1093/brain/109.2.357
- Sperry, R. W. (1982). Some effects of disconnecting the cerebral hemispheres. *Science, 217*(4566), 1223–1226. https://doi.org/10.1126/science.7112125
- Tononi, G., Boly, M., Massimini, M., et al. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience, 17*, 450–461. https://doi.org/10.1038/nrn.2016.44.
- Volz, L. J., & Gazzaniga, M. S. (2017). Interaction in isolation: 50 years of insights from split-brain research. *Brain, 140*(7), 2051-2060.
- Watzl, S. (2014). Attentional organization and the unity of consciousness. *Journal of Consciousness Studies, 21*(7–8), 56–87. http://www.ingentaconnect.com/content/imp/jcs/2014/00000021/F0020007/art00004
- Watzl, S. (2017). *Structuring Mind*. Oxford: OUP.
- Wiese, W. (2022). Attentional structure and phenomenal unity. *Open Philosophy, 5*(1), 254–264. https://doi.org/10.1515/opphil-2022-0197

Elizabeth Schechter (University of Maryland, College Park): "Self-Consciousness after Split-Brain Surgery"

In this talk, I first argue that the two hemispheres of a split-brain subject are associated with distinct conscious thinkers and, indeed, distinct thinkers of self-conscious thoughts, R and L. I then argue that the dynamics of self-conscious thought after split-brain surgery shows that R and L are unlike other pairs of self-conscious thinkers. Indeed, there is a basic psychological capacity much simpler than but ordinarily inherent in full-blown psychological self-consciousness that R and L both lack.