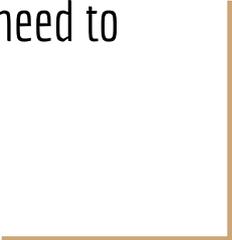




Octopuses, Split-Brains, and the Universe

How unified do subjects need to
be?



Goal of this talk

I want to bring various sorts of empirical evidence into conversation with a metaphysical doctrine which is recurrently popular but seems to be obviously false.

Specifically, I'm interested in whether cases where conscious unity is impaired or partial can affect arguments against 'subject-monism'.

Plan of this talk

1. A question: how unified do subjects need to be?
2. Four answers to the question
3. The Obvious Argument against monism
4. Some interesting empirical cases:
 - a. Split-Brain Patients
 - b. Octopuses
5. What do the empirical cases mean for the four answers?
6. Implications for the Obvious Argument

Section 1: The Question

What connects the split-brain, octopuses, and subject-monism can be captured as a specific question:

How far does the unity or dis-unity among a set of experiences constrain whether they can be ascribed to a single subject?

That is, does knowing that certain experiences occur, and how they are related, tell us who they belong to?

This question demands two preliminary clarifications...

Section 1: The Question

First clarification: what does 'unity' (or 'dis-unity') mean?

Unity is the various ways that experiences of the same person tend to 'hang together' (cf. Schechter 2010, Bayne & Chalmers 2003, Tye 2003).

We can identify several particular forms of unity:

- Global coherence
- Representational composition
- Causal interdependence
- Access-unity
 - Behavioural unity
 - Inferential unity
 - Meta-representational unity

Section 1: The Question

Global coherence: When the contents of multiple experiences are consistent with one another, with any conflict tending to produce experienced tension.

(E.g. when I feel good about an event and you feel bad about the same event, our experiences are not globally coherent, because there need be no tension that results from their conflict.)

Representational composition: When the contents of multiple experiences are subsumed by a more complex content belonging to a composite experience.

(E.g. when I experience redness and squareness, I can experience them as aspects of a single red square; when I experience redness and blueness, I can experience them as two contrasting colours.)

Causal interdependence: When one experience happens the way it does as a direct result of how another another experience is happening.

(E.g. if you changed one of my visual experiences, that would have direct knock-on effects on other aspects of my stream of consciousness.)

Section 1: The Question

Access-unity: When the content of multiple experiences, *and* the conjunction of their contents, are available for the same cognitive uses.

Behavioural unity: When the same behavioural systems are guided by the contents of multiple experiences.

(E.g. my evasive behaviour is prompted by my pain but not by yours, and vice versa.)

Inferential unity: When the contents of multiple experiences are subsumed by a more complex content belonging to a composite experience.

(E.g. if I represent that P and that $P \rightarrow Q$, I will usually conclude that Q ; if I represent that P and you represent that $P \rightarrow Q$, neither of us is likely to conclude that Q .)

Meta-representational unity: When verbal reports, introspective thoughts, or later memories record both experiences together.

(E.g. I can report my current perceptions, but not yours, and vice versa.)

Section 1: The Question

Second clarification: what does 'subject' mean?

The way we talk about consciousness involves both *someone* who is undergoing something, and *something* that is happening in or to them.

The term 'subject' is meant to pick out the former of these (which we could also call a person, soul, self, or possibly a mind).

We might define a subject of experience as an entity which it is like something to be. If I ask what it is like to be this dog, or this chair, or something else, I am thereby presupposing that the dog, chair, or whatever, is a subject.

Plausibly, subjects are what we mean when we use personal pronouns, and quantifiers like 'someone'.

Section 2: Four Answers

How far does the unity or dis-unity among a set of experiences constrain whether they can be ascribed to a single subject?

Compare four possible answers:

- A. Significant constraint, because of a causal connection (belonging to the same subject causally guarantees a high degree of unity);
- B. Significant constraint, because of a constitutive connection (belonging to the same subject *just is* having a high degree of unity);
- C. No significant constraint, because subjects are not defined by their mental lives but by their underlying substrate (e.g. biological boundaries of the organism);
- D. No significant constraint. Even if subjects are defined by their mental lives, those mental lives can potentially incorporate even radical disunity.

Section 2: Four Answers

The idea of answer **A** is that while belonging to the same subject and being unified are different things, the one is in fact a guarantee of the other.

'If a single subject has two experiences, those experiences couldn't help but be unified.'

Answer **B** agrees that subjects only have unified experiences, but explains this as a conceptual fact about what we mean by 'subject'.

'Dis-unified experiences can't be counted as belonging to the same subject - that's the point of talking about subjects.'

Section 2: Four Answers

Answer **C** expresses the thought that the identity and boundaries of a conscious thing need not depend on anything about its consciousness.

'Some beings, whether organisms, organs, souls, material bodies, etc., happen to be conscious - but that doesn't affect their individuation. (Compare: the identity and boundaries of a dancer don't depend on its dancing.)'

Answer **D** is the rejection of the other three answers: there is a distinctly experiential fact about whether two experiences belong to the same subject, but it's dissociable from unity relations.

Section 3: The Obvious Argument Against Monism

There are many reasons to be interested in the question:

How far does the unity or dis-unity among a set of experiences constrain whether they can be ascribed to a single subject?

One is that it's relevant to how we should understand various unusual cases (e.g. dissociative identities) or hypotheticals (e.g. technologically-enabled telepathy).

Another is that it's relevant to how we should evaluate our everyday assumptions about the number and identities of subjects.

Section 3: The Obvious Argument Against Monism

The everyday picture: There are at least 7 billion distinct subjects of experience on earth, each corresponding to a separate human body.

This picture is baked into our everyday ways of talking, acting, and thinking.

(Even if we accept occasional exceptions, e.g. two subjects in one body for certain conjoined twins)

To throw it into sharper focus, let's contrast it with a rival picture, 'subject-monism'...

Section 3: The Obvious Argument Against Monism

The monistic picture: There is 1 subject of experience in the universe, undergoing the experiences associated with multiple human (and other) bodies.

The monistic picture can be found in a lot of places:

- Contemporary philosophers (sometimes described as ‘cosmopsychism’ - see e.g. Shani 2015, Goff Forthcoming)
- Historical philosophers (with varying degrees of qualification, we might mention Spinoza, Schopenhauer, and Hegel)
- At least one major religious tradition (the Vedanta versions of Hinduism)
- Various mystical traditions within other religions

Section 3: The Obvious Argument Against Monism

Distinguish three slightly different versions of subject-monism:

- **Reductive:** Distinct subjects like you and I do exist, but only as derivative fragments of the one mind;
- **Eliminative:** Distinct subjects like 'you' and 'I' do not exist - there is only the one mind;
- **Identifying:** The subjects we refer to as 'you' and 'I' do exist, but are not distinct - I am the one mind, and so are you.

(It's an open question how different these three really are - they may just be different ways of saying the same thing.)

Section 3: The Obvious Argument Against Monism

At first glance, the monistic picture is clearly false. There is decisive evidence against it, and in favour of the everyday picture.

Namely, the obvious dis-unity among the various human experiences on earth. Call this 'the Obvious Argument against monism':

Premise 1: Experiences that are dis-unified cannot belong to the same subject.

Premise 2: Many experiences on earth are dis-unified.

Conclusion: The experiences on earth do not all belong to the same subject.

Section 3: The Obvious Argument Against Monism

The Obvious Argument is so obvious that it's not often presented explicitly. But here's Philip Goff making it in 2012:

“At this moment, I am experiencing no pleasure, but I do feel a terrible pain in my knee... Sitting across from me is Dave, who is eating an ice cream. He tells me he is currently enjoying a pain free existence, and is in fact extracting great pleasure from his ice cream... It follows that there are currently in existence two things: one that feels pain but no pleasure, and another that feels pleasure but no pain. Therefore, existence monism is false.”

(Goff, 'There Is More Than One Thing', 2012, p.113)

Section 3: The Obvious Argument Against Monism

The Obvious Argument is not at all new: here's Ibn Sina making it in 1027:

“And we know that the soul is not one in all the bodies. If it were one and in relation many, it would have knowledge or ignorance in all and it would not be hidden for Zayd what is in the soul of Amr... Therefore the soul is not one.”

(Ibn Sina/Avicenna, 'The Book of Healing', 1027, Book 5, Chapter 3, Sections 7-8)

Section 3: The Obvious Argument Against Monism

But clearly, premise 1 of the Obvious Argument ('Experiences that are dis-unified cannot belong to the same subject') depends on how we answer the question we started with...

Answers **A** and **B** support the Obvious Argument, and thus the everyday world-picture. They make observed disunity among experiences into good evidence for their belonging to multiple subjects.

Answers **C** and **D** undermine premise 1, and thus undermine the Obvious Argument for the everyday picture. They suggest that experiences may belong to the same subject even if they are dis-unified.

Section 4: Cases

Let's look at two empirical cases where different sorts of unity come apart.

Both cases have been discussed by philosophers as challenges for a theory of subjectivity.

It's sometimes suggested that in light of cases like this, it may be harder to maintain answers A or B in strong forms. I agree, and I think this weakens the Obvious Argument, and indirectly helps monism.

4.a The Split-Brain Phenomenon

The 'split-brain' phenomenon results from the surgical severing of the corpus callosum, the bundle of nerve fibres that connects the two hemispheres of the human brain.

When this is cut (which was done to cure serious forms of epilepsy) the two hemispheres lose one major means of communication.

(NB: They do not lose all means of communication: the brainstem, including the thalamus, remains uncut, and constitutes a major connection.)

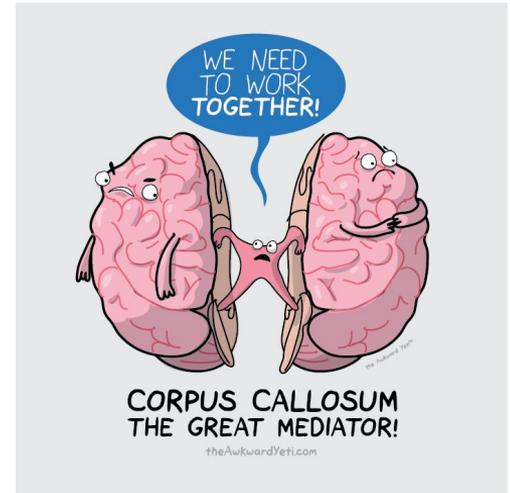


Image credit: Nick Seluk,
<http://theawkwardyeti.com/comic/yay-brain-anatomy/>

4.a The Split-Brain Phenomenon

Under normal conditions, most of the time, the patients appear pretty much normal: you wouldn't know by talking to them that their brain is 'split'.

But very strange results can be elicited by carefully segregating inputs to their hemispheres - typically by presenting images to only one half of the visual field, without allowing the eyes to move and bring it into the other half.

This ensures that sensory information goes to one hemisphere but not the other. The result is that the patient can respond to that information only using the motor organs and abilities associated with that hemisphere.

4.a The Split-Brain Phenomenon

When an image is presented to the right visual field, and processed by the left hemisphere, the patient can report what they saw, but can only pick it out using their right hand.

When an image is presented to the left visual field, and processed by the right hemisphere, the patient reports seeing nothing - but can pick out the object they saw if they use their left hand.

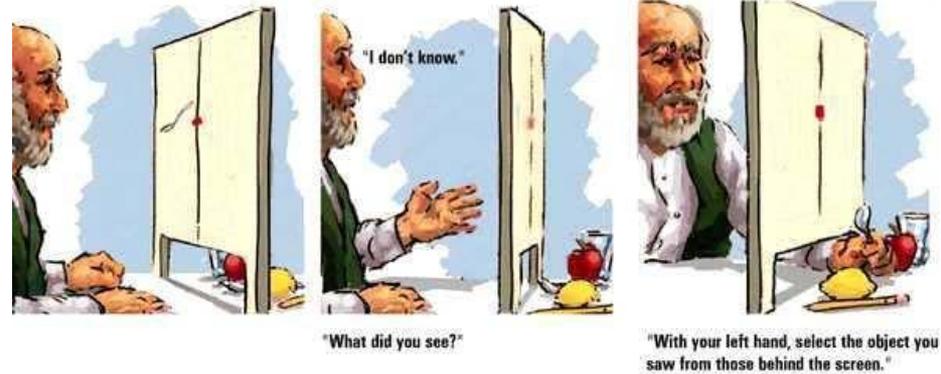


Image credit: Brodie Sprlyan <https://futurethoughtblog.wordpress.com/2016/04/19/of-two-minds/>

4.a The Split-Brain Phenomenon

This seems to show that perceptual experiences based in the two hemispheres are access-dis-unified. In particular, they are:

- Not meta-representationally unified (they cannot be reported together)
- Not inferentially unified (they cannot be used together in reasoning)
- Not behaviourally unified (they cannot be used to guide the same activities)

There is also a lack of representational composition: if shown 'toad' and 'stool' in different visual fields, the patient will draw a toad (with one hand) or a stool (with the other), but not a toadstool.

4.a The Split-Brain Phenomenon

This could be taken to show that the two sets of experiences belong to two different subjects - a 'left-hemisphere subject' and a 'right-hemisphere subject'.

However, there are also important sorts of unity among these experiences that seem to tell in favour of seeing them as a single subject.

For one thing, there is overwhelming global coherence in their values, beliefs, memories, actions, etc. They maintain a consistent persona and life-plan.

4.a The Split-Brain Phenomenon

This presumably depends on (thalamic?) mechanisms that establish causal interdependence, so that the two hemispheres don't start developing separate ideas about themselves and their goals.

Moreover, there is representational composition with some of their experiences, e.g. one could detect contrast or matching between shapes felt with both hands, but not between a hand and the contralateral visual field (Zaidel 1998).

And there is the behavioural unity they display in normal tasks (including following the experimenter's instructions to walk around, sit down, stare at a screen, etc.).

4.a The Split-Brain Phenomenon

In light of this, it's hard to know whether the experiences associated with the two hemispheres belong to the same subject or not.

Some philosophers have even claimed that we can't meaningfully describe matters in either of those two ways.

E.g. Nagel 1971, p.409:

"I do not believe it is possible to accept... any conclusion involving the ascription to them of a whole number of minds."

4.b The Octopus Brain(s)

Octopuses (and other cephalopods: squid, cuttlefish, nautilus) come from a lineage that diverged from our ancestors at least 500 million years ago: our common ancestor back then almost certainly did not have a large brain.

This means octopuses evolved Intelligence and large brains independently of vertebrates.

As a result, their nervous system is organised very differently from ours and those of our relatives.



Image from Edel, Baars, & Seth 2005, p.178

4.c The Octopus Brain(s)

Whereas most of our neurons are concentrated in one place (the brain), two-thirds of the octopus's neurons are found in its arms.

This led researchers to wonder if the arms were guiding their activities independently of the central system.

To test this, the arms were either removed from the octopus entirely, or had their nervous connections to the central system severed (Sumbre et al. 2001, Gutfreund et al. 2006).

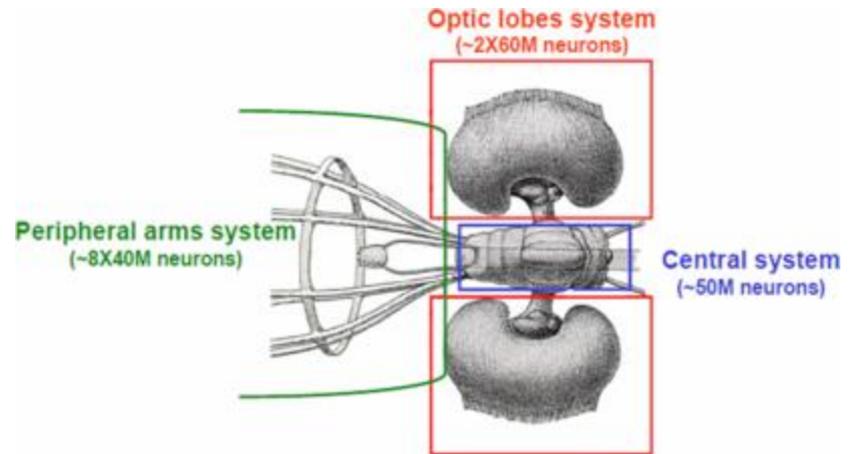


Image from Kang et al., 2013, p.3

4.c The Octopus Brain(s)

The result was that the arm's behaviour was largely unaffected(!): in particular, the arms continued to...

- Crawl around
- Reach out for things that touch them
- Withdraw from painful acids
- Detect food and try to pass it 'down' to where the mouth would be

This looks like a major lack of causal interdependence between whatever experiences (if any) are associated with the arms and whichever experiences (if any) are associated with the central system. What happens to the one has only modest direct effect on the other.

4.c The Octopus Brain(s)

Indeed, octopuses often have trouble with tasks requiring them to move a single specific arm.

(It almost seems as though the octopus brain is travels around with 8 over-eager helpers who all do their best to follow its shouted instructions.)

However, experiments by Gutnick et al. (2011) have succeeded in showing that an octopus can use visual information (available only to the central system) to guide a single arm's movements through a maze.

4.c The Octopus Brain(s)

This seems to show a significant sort of access-unity among processes in the arm and the brain, though one that takes a while to engage.

Should we conclude that the octopus is:

- One subject, aware only of what's happening in the brain;
- One subject, aware in a dis-unified way of what's happening all through the nervous system;
- Nine subjects, one smarter than the rest;
- Not conscious at all;
- Something else? (Godfrey-Smith, 2016, p.103, writes: *"The octopus may be in a sort of hybrid situation. For an octopus, its arms are partly self—they can be directed and used to manipulate things. But from the central brain's perspective, they are partly non-self too, partly agents of their own."* cf. Carls-Diamante 2017)

Section 5: Implications for the Four Answers

These cases show various sorts of dissociation between different types and degrees of unity. What does this mean for the four answers to our starting question?

How far does the unity or dis-unity among a set of experiences constrain whether they can be ascribed to a single subject?

- A. Significant constraint, because of a causal connection;
- B. Significant constraint, because of a constitutive connection;
- C. No significant constraint, because subjects are not defined by their mental lives but by their underlying substrate;
- D. No significant constraint, because subjects are defined by their mental lives but those lives can potentially incorporate even radical disunity.

Section 5: Implications for the Four Answers

These cases don't say much about answers C and D, because those answers explicitly allow that unity doesn't say much about subjecthood.

But they seem to be evidence against answer A. If belonging to the same subject is the single factor that explains all the forms of unity, why do we see those forms dissociating? Wouldn't we expect to find them always going together?

Answer A seems to say that consciousness always comes in 'units': we seem to have found cases where it comes in 'messier', more open-ended, sorts of system.

Section 5: Implications for the Four Answers

The evidence against answer A isn't decisive: it's possible that belonging to the same subject really does guarantee unity, but that other things can also produce unity.

E.g. maybe the octopus is 9 subjects, not 1, but the forms of behavioural unity it displays are accomplished by between-subjects sorts of co-ordination.

So answer A isn't ruled out. But we might find our confidence in it significantly shaken: if the world is full of such strange and unexpected sorts of mind, why think we know in advance what sorts of dis-unity are and are not possible for a single subject?

Section 5: Implications for the Four Answers

What about answer B?

If 'belonging to the same subject' just means 'unified', then it starts to look like the various experiences in these cases neither do nor don't belong to the same subject.

(Or both do and don't belong to the same subject.)

The very concept of a subject is 'stretched to breaking point' (recall Nagel and Godfrey-Smith).

Section 5: Implications for the Four Answers

A natural move here is to say that rather than *abandoning* the concept of a subject defined by unity, we should just *complicate* it.

Maybe belonging to the same subject is a matter of degree: different things can count as a single subject relative to different ways of 'calibrating' the concept.

(E.g. if we weight representational composition highly, the hemisphere-systems of a split-brain patient are two subjects; if we weight it less, they count as parts of a single subject.)

Section 5: Implications for the Four Answers

This seems to be a reasonable thing to do with other concepts, e.g. the concept 'organism'.

Are my cells organisms? Is an ant colony an organism? Are the different cells in a slime mould or a sponge parts of the same organism?

It seems perfectly appropriate to say 'on one way to calibrate the concept 'organism', they are: on another way, they aren't'. There's no need to determine the single right definition.

Section 6: Implications for Subject-Monism

So:

The cases considered so far seem to be (mild) evidence against answer A, and to force a 'relativisation' of answer B.

Answers A and B were also the two that strongly supported the Obvious Argument against subject-monism.

Thus, these cases indirectly help subject-monism. This is fairly easy to see with answer A: if it's false, then it can't be used to support the Obvious Argument.

But things are more complicated with answer B.

Section 6: Implications for Subject-Monism

If there are many equally good ways to calibrate the concept 'subject', what does subject-monism even mean?

I suggest a charitable reading: subject-monism should be construed as saying that there is *at least one* interesting way to calibrate the concept 'subject', on which all experiences belong to one subject.

This allows for there to be other ways to calibrate the concept, which would vindicate the everyday picture. The two pictures may be compatible.

Section 6: Implications for Subject-Monism

What sort of calibration would yield that result?

Firstly, one might focus on direct causal interdependence, and relax any 'strength' requirement' for the direct causal effects of one experience on another.

Since any two brain-events exert some direct force on one another, even if very distant, they are all causally interdependent.

Secondly, one might focus on representational composition, and relax any 'likelihood requirement' for the composition to actually occur.

Since there's no in-principle barrier to building futuristic physical connections between my brain and yours, our experiences are *capable* of being compounded, even if the necessary events are very unlikely to happen.

Section 6: Implications for Subject-Monism

A worry: if subject-monism is based on this extreme way of calibrating the concept 'subject', doesn't it become entirely trivial?

(*Of course* there's only one subject, if that just means that all experiences can interact, and could in principle be compounded with each other!)

But subject-monism in this sense isn't strictly trivial: it might have been false.

In a 'Cartesian' world, filled with indivisible, immaterial souls, there might not be any *direct* interaction between experiences in different souls, nor even the *in-principle* possibility of representational composition between them.

Section 6: Implications for Subject-Monism

So subject-monism is at least expressing the denial of a Cartesian dualist world.

A further worry: Surely this watered-down version of the concept 'subject' is of basically no scientific use. Why should we take any interest in it?

Two quick replies:

1. A reason not to bother saying something isn't a reason to deny that it's true.
2. The concept 'subject' isn't just a scientific concept: it also serves to organise our sympathies and prudential or moral concerns. The subject-monist thesis may be of interest for its impact on this.

Section 6: Implications for Subject-Monism

So where does all this leave us?

I considered four answers to the question:

How far does the unity or dis-unity among a set of experiences constrain whether they can be ascribed to a single subject?

- A. Significant constraint, because of a causal connection;
- B. Significant constraint, because of a constitutive connection;
- C. No significant constraint, because subjects are not defined by their mental lives but by their underlying substrate;
- D. No significant constraint, because subjects are defined by their mental lives but those lives can potentially incorporate even radical disunity.

Section 6: Implications for Subject-Monism

Answers A and B directly support the Obvious Argument against subject-monism.

Consideration of octopuses and split-brain patients makes it harder to maintain answer A.

They also push us towards a relativised version of answer B, which may end up making subject-monism compatible with the everyday picture.

Answers C and D do not directly bear on the subject-monist position.

Cautious result: Subject-monism is not as obviously false as it seems.

References:

- Bayne, Timothy. 2008. 'The Unity of Consciousness and the Split-Brain Syndrome.' *The Journal of Philosophy* 105(6): 277-300.
- Bayne, Timothy. 2010. *The Unity of Consciousness*. Oxford: Oxford University Press
- Bayne, Timothy, & Chalmers, David. 2003. 'What is the Unity of Consciousness?' In Cleeremans, A. (ed.), *The Unity of Consciousness: Binding, Integration, Dissociation*, Oxford: Oxford University Press: 23-58
- Carls-Diamante, Sidney. (2017). "The octopus and the unity of consciousness." *Biology and Philosophy* 32 (6): 1269-1287.
- Edelman, David, Baars, Bernard, and Seth, Anil. (2005). "Criteria for consciousness in humans and other mammals." *Consciousness and Cognition* 14: 119-139.
- Gazzaniga, Michael, Bogen, Joseph, and Sperry, Roger. 1962. 'Some Functional Effects of Sectioning the Cerebral Commissures in Man.' *Proceedings of the National Academy of Sciences* 48(2): 17-65
- Godfrey-Smith, Peter (2016). "Other Minds: The Octopus, the Sea, and the Deep Origins of Consciousness." MacMillan
- Goff, P. (2012). *There is More than One Thing*. In Goff, P. (ed) *Spinoza on Monism*. Palgrave MacMillan, pp. 113-122
- Goff, P. (forthcoming). *Cosmopsychism, Micropsychism, and the Grounding Relation*. In William Seager (ed.), *Routledge Panpsychism Handbook*. Routledge.
- Gutfreund, Y., Henry Matzner, Tamar Flash, and Binyamin Hochner (2006) *Patterns of Motor Activity in the Isolated Nerve Cord of the Octopus Arm* *The Biological Bulletin* 211: 212-222.
- Gutnick T, Byrne RA, Hochner B, Kuba M (2011) *Octopus vulgaris* uses visual information to determine the location of its arm. *Current Biology*. 21(6):460-462.
- Ibn Sina/Avicenna (1952). *The Book of Healing*. Originally published 1027, in Rahman, F. (trans.) *Avicenna's Psychology*. Oxford, Oxford University Press.
- Kang, Rongjie, Branson, David, Zheng, Tianjiang, Guglielmino, Emanuele, and Caldwell, Darwin. (2013). "Design, modeling and control of a pneumatically actuated manipulator inspired by biological continuum structures." *Bioinspiration & Biomimetics* 8(3): 1-14.
- Marks, Charles. 1980. *Commissurotomy, Consciousness, and Unity of Mind*. MIT Press
- Nagel, Thomas. 1971. 'Brain Bisection and the Unity of Consciousness.' *Synthese* 22:396-413.
- Schechter, E. (2009). *Persons and Psychological Frameworks: A Critique of Tye*. *Journal of Consciousness Studies* 16(2-3): 141-163
- Schechter, E. (2010). *Two Unities of Consciousness*. *European Journal of Philosophy*, V21:2, pp.1-25
- Shani, I. (2015). *Cosmopsychism: A Holistic Approach to the Metaphysics of Experience*. *Philosophical Papers* 44 (3):389-437.
- Sperry, Roger. 1964. 'Brain Bisection and Mechanisms of Consciousness.' In Eccles, John (ed.), *Brain and Conscious Experience*. Springer-Verlag
- Sumbre, G., Gutfreund, Y., Fiorito, G., Flash, T., & Hochner, B. (2001). *Control of octopus arm extension by a peripheral motor program*. *Science*, 293, 1845-1848.
- Tye, Michael. 2003. *Consciousness and Persons: Unity and Identity*. MIT Press
- E. Zaidel, "Stereognosis in the Chronic Split-Brain: Hemispheric Differences, Ipsilateral Control and Sensory Integration Across the Midline," *Neuropsychologia* 36/11 (1998): 1033-47