

The Unity of Consciousness: A Cartography

Tim Bayne
University of Oxford and St. Catherine's College
Manor Road
Oxford OX1 3UJ
United Kingdom
tim.bayne@gmail.com

This paper was published in (2007) M. Marraffa, M. de Caro & F. Ferretti (eds.) *Cartographies of the Mind: Philosophy and Psychology in Intersection*. Dordrecht: Kluwer, 201-10. Please consult the published version for purposes of quotation.

1. Introduction

One of the many fault-lines within accounts of consciousness concerns the unity of consciousness. Some theorists claim that consciousness is unified—indeed, some theorists insist that consciousness is essentially unified. Other theorists assert that the unity of consciousness is an illusion, and that consciousness is often, if not invariably, disunified. Unfortunately, it is rare for proponents of either side of the debate to explain what the unity of consciousness might involve. What would it mean for consciousness to be unified? In this chapter I provide a brief cartography of the unity of consciousness. In the next section I introduce a number of unity relations that can hold between conscious states, and in the following sections I show how these unity relations can be used to construct various conceptions of the unity of consciousness—what I call *unity theses*. These unity theses provide us with a set of reference points by means of which we can orient discussions of the (dis)unity of consciousness.

2. Unity relations

A number of unity relations structure consciousness. Here I will introduce four such relations.

The first unity relation is *subject unity*. Conscious states are subject unified when they are had by the same subject of experience. My current experiences are mutually subject unified, and your current experiences are mutually subject unified.

A second unity relation—or rather, type of unity relation—is *representational unity*. Representational unity concerns the representational content of experience. One form of representational unity involves the integration of representational content based around perceptual objects. How does the brain put together representations of colour, shape, motion, and category to generate a unified percept of (say) a tiger stalking its prey? This is a central version of the binding problem.

A third unity relation is *access unity*. The contents of experience are often available to ‘consuming systems’—the kinds of systems responsible for intentional behavior, verbal report, belief-formation, perceptual categorization, memory consolidation, the voluntary allocation of attention, and so on. Roughly speaking, the members of a set of experiences will be access unified to the extent that their contents are available to the subject’s consuming systems.

A final unity relation is *phenomenal unity*. Experiences are phenomenally unified when they have a conjoint phenomenology; that is, when there is something it is like to experience them together. There is something it is like to have an experience of pain, there is something it is like to see a dog, and there is something it is like to have an experience of a dog and an experience of pain together (Bayne & Chalmers 2003). One can think of phenomenal unity as a relation that experiences have when they occur as components of a single phenomenal state.¹

¹ In the recent philosophical literature phenomenal unity is frequently referred to by the term ‘co-consciousness’ (Dainton 2000; Lockwood 1989; Hurley 1998; Shoemaker 2003). I prefer ‘phenomenal unity’ rather than ‘co-consciousness’ because ‘co-consciousness’ has been used within psychology for a very different relation (roughly, the relation that experiences bear to each other when they are co-subjective but not phenomenally unified) and because ‘phenomenal unity’ is more descriptively accurate than ‘co-consciousness’.

With these unity relations in hand we are now in a position to examine various conceptions of the unity of consciousness.

3. The consistency thesis

One conception of the unity of consciousness focuses on representational consistency. According to the *consistency thesis*, phenomenally unified experiences must be representationally consistent. Baars appears to have the consistency thesis in mind when he suggests that the unity of consciousness involves the idea that “the flow of conscious experience [...] is limited to a single internally consistent content at any given moment” (Baars 1993, 285). Baars himself holds that consciousness is unified in this sense, and a number of other theorists have agreed with him (e.g. Tononi & Edelman 1998).

But is the consistency thesis correct? As the phenomenon of binocular rivalry shows, consciousness is clearly resistant to representational inconsistency. When one’s eyes are presented with different images—say, a sunflower and a house—one’s visual experience typically alternates between these two images. One typically sees either the house or the sunflower but not both at the same time (Rubin 2003). Nonetheless, imaging experiments reveal that both stimuli are processed (Logothetis et al. 2003). But if both images are perceptually processed, why are we only conscious of one at a time? Perhaps the mechanisms of consciousness ensure that inconsistent percepts cannot be phenomenally unified. The phenomenon of multi-stable images points to a similar conclusion. One can see a Necker cube as having one or other of two orientations, but one cannot see it as having both orientations at once. (Try it!)

But although consciousness resists representational inconsistency the consistency thesis appears to be false. Consider the effects of wearing inverting spectacles (Kohler 1961; Taylor 1962). Prior to adaptation, such spectacles cause the contents of vision to be inconsistent with the contents of touch: the vase looks as though it’s upside down, but it feels as though it’s the right way up. Of course, the data from these experiments shows that consciousness exhibits a *drive* towards consistency, for the contents of one’s various sensory modalities are, after a

period, brought into line with each other. But my point here is that consciousness tolerates inter-modal perceptual inconsistency, if only temporarily.

One might challenge the claim that these cases involve inconsistent contents.² Perhaps the fact that that the visual and tactile experiences involve different modalities entails (or at least suggests) that their contents differ. Perhaps we should not think of the experiences in different perceptual modalities as having the same intentional objects. But I am inclined to resist this objection. It seems to me that the representational contents of the two experiences are inconsistent, in that they could not be simultaneously satisfied. (Some support for this can be derived from the fact that the two experiences justify inconsistent beliefs.)

One might attempt to revise the consistency thesis so as to deal with inter-modal inconsistency. The obvious way to do this is to restrict the consistency constraint to intramodal contexts (Hurley 2000). According to the revised consistency thesis, phenomenally unified experiences within a single modality must be representationally consistent. (I will leave to one side the tricky question of how perceptual modalities might be individuated.) Does the revised consistency thesis fare any better than the original version? Perhaps not.

Consider the waterfall illusion, first described by Aristotle. If one looks at a waterfall for a period of time, and then directs one's gaze to a stationary object, the stationary object will appear to be moving in the direction opposite to the apparent motion of the first object. Yet, at the same time, it also appears to be stationary.³ Here is Frisby's description of the experience produced by the waterfall illusion:

Although the after-effect gives a very clear illusion of movement, the apparently moving features nevertheless seem to stay still! That is, we are still aware of features remaining in their 'proper' locations even though

² Thanks to Chris Maloney for the following objection.

³ See Frisby (1979, 100-101) for a description of how to produce the waterfall illusion using a record player. Those readers born after 1980 might have to borrow their parents' record player.

they are seen as moving. What we see is logically impossible! (Frisby 1979: 101)

Are we simultaneously conscious of the target object as both moving *and* not-moving? Frisby seems to think so, and a number of authors have joined him in this (see Crane 1988). I'm not so sure. I can see the object as moving, and I can see it as stationary—and I can move my attention between these two experiences—but I am not convinced that I ever experience it in both ways at once.⁴

So-called impossible objects provide another potential counter-example to the revised consistency constraint (Tye 2003). One of my favourites is the devil's pitchfork, also known as the impossible fork:




Figure 1: The Devil's Pitchfork

Does your experience of the Devil's Pitchfork have inconsistent contents? It might be thought so—after all, how else could one see it as an impossible object? But perhaps one doesn't see it as an impossible object *all at once*. Perhaps one only builds up a perception of it as an impossible object over time. Perhaps there is no point at which one's visual phenomenology has inconsistent content. When one focuses on the prongs one experiences them as straight, and when one focuses on the 'handle' one experiences the middle prong as lying behind the upper and lower prongs, but one does not experience both ends of the pitchfork

⁴ Another possibility is that one sees the target object as moving-relative-to-object-X and as not-moving-relative-to-object-Y.

at once. One builds up a representation of the pitchfork as a whole; this representation has impossible content, but it is not a visual representation. It is, rather, a belief formed on the basis of visual experience. (It is a belief about the kind of object that would have to exist in order for one's experiences of the various parts of the object to be veridical.) I am not sure that representations of so-called impossible objects falsify the revised consistency constraint, but I think it does put pressure on it.

Suppose that the revised (intra-modal) consistency thesis is true—what might follow from this? The original consistency thesis suggested that something in the very nature of consciousness might prevent inconsistent contents from occurring within a single state of consciousness. By contrast, the revised consistency thesis promises to tell us more about perceptual modalities and representational content than about consciousness as such.

4. The availability thesis

Another common conception of the unity of consciousness focuses on access unity. According to the *availability thesis*, the contents of a subject's conscious states are globally available—they are available to each of the consuming systems that the subject possesses at the time in question. If the availability thesis is correct, then the contents of a subject's current experiences will be available to the same consuming systems, and hence their consciousness will be behaviourally unified.

Although it is rarely put in quite these terms, something very much like the availability thesis is widely assumed within current approaches to consciousness. Consider, for example, global workspace models of consciousness (Baars 1988, 2002; Dehaene & Naccache 2001; Dennett 2001). According to such models, conscious content occurs within a workspace that enables it to be broadcast throughout the subject's cognitive system. Sometimes global workspace theorists appear to *identify* consciousness with global availability, while in other places they appear to suggest that consciousness is the categorical ground of global availability. Either way, global workspace models hold that the difference between conscious and unconscious content consists in the range of consuming

systems to which the content is available: unconscious content is available only to a restricted range of consuming systems, whereas conscious content is globally available to the subject's consuming systems.

The availability thesis accords with normal waking consciousness, but it fits less well with so-called altered states of consciousness, such as dreaming and delirium. The strangest events occur in dreams—a dog turns into an elephant, one's aged grandmother eats a hamster on toast, and the Queen gets married to Lenny Bruce—without the dreamer registering any awareness of the incongruity. Although the contents of dream consciousness may be available to some of the subject's consuming systems, they do not seem to be available to those systems involved in consistency checking, belief-updating, introspection, and (in at least some cases) memory consolidation. The tight correlation between consciousness and global availability seems also to be compromised in delirium. Delirious subjects lack the integrative, monitoring and mnemonic capacities typically associated with consciousness (Gill & Mayou 2000; Lipowski 1990; Fleminger 2002). In short, the association between global availability and waking consciousness appears to have more to do with the fact that the subject is awake than with the fact that the subject is conscious.

There are at least two ways in which the proponent of the availability thesis might respond to these objections. Firstly, she might hold that consciousness can come in degrees, and that the restricted availability of content that one sees in dreams and delirium is a reflection of the fact that these states involve only minimal degrees of consciousness. I do not have much sympathy with this response—whatever exactly a 'minimal' form of consciousness involves, it seems to me that there is little reason to think that the phrase applies to dreams and delirium. It seems to me that the phenomenology involved in dreams and delirium is—or at least can be—as rich as that which occurs in normal waking consciousness. Secondly, the proponent of the availability thesis might hold that the correlation between consciousness and global availability holds only in the normal waking state. This is a more plausible response, but in making it the theorist has given up on a global availability approach to consciousness as such.

But the global availability account of consciousness is problematic even if we restrict our attention to the normal waking state. In the remainder of this section I examine a number of syndromes, each of which suggest that consuming systems can have *differential* access to the contents of consciousness: content can be available to some of the subject's consuming systems without being available to all of their consuming systems.

In the Dimensional Change Card Sort task children are shown two target cards and asked to sort a series of cards (e.g. red rabbits and blue cars) into piles according to a certain dimension such as colour (Zelazo 1996, 2004). Having sorted several cards, children are then told to switch the sorting rule. 3 year-olds typically fail to switch dimensions regardless of which dimension is presented first. Nonetheless, the children usually respond correctly to questions about what they ought to be doing. Their verbal behaviour suggests that they are conscious of the post-switch rules, yet their sorting behavior suggests that they are conscious of only the pre-switch rules. In short, the children exhibit a dissociation between knowing and doing.

Research in the visual perception of normal subjects provides other examples of selective availability to consuming systems. Consider the following experiment conducted by Cumming (reported in Allport 1988). Cumming's subjects were shown a horizontal row of five letters flashed in rapid succession, one after another. This spatio-temporal arrangement of the letters was designed to produce a form of metacontrast masking known as 'sequential blanking'. The subjects were then given a visual search task. The instructions were to press one key if the letter J (for example) was present in the display, and to press another key if no J was present.

When urged to respond as fast as possible, even at the cost of making a good many errors, subjects now tended to respond to the occurrence of a target letter in the 'blanked' positions with a fast (and *correct*) press of the 'target present' key, and then, a moment later, to apologize for having made an error. (Allport 1988, 175)

Arguably, Cumming's subjects had experiences whose contents were available for report in one modality (manual button-pressing) but not another (verbal report).⁵

Marcel's experiments provide examples of similar dissociations (Marcel 1993, 1994). Marcel asked normal subjects to respond when they detected a light that was illuminated for 200 ms. Subjects were asked to respond as quickly as possible, and to respond in three ways at once: by blinking, by pressing a button, and by saying 'yes'. Subjects often gave inconsistent responses. In a single trial a subject might, say, report with their finger that he saw the light but fail to report orally that he saw a light. Or a subject might indicate that she had seen the light by pressing the button but fail to say 'yes' or blink. Marcel also found that subjects were unaware that their responses were inconsistent.⁶ Again, subjects appear to have had conscious states whose contents were not globally available to consuming systems. On any particular trial, an experience of the light might (say) be available to those consuming systems involved in button-pressing but not verbal report, or vice-versa.

The clinical syndrome known as 'anosognosia' also provides prima facie evidence against the availability thesis. Anosognosia involves a lack of awareness of a deficit—typically unilateral neglect or hemiplegia (paralysis on one side of the body). Often this lack of awareness is partial, and the patient will have 'dim knowledge' of their condition (Bisiach & Berti 1995; Marcel et al. 2004; Vuilleumier 2004). A patient with hemiplegia may verbally acknowledge her condition but nonetheless attempt to rise from bed or engage in other activities

⁵ It is possible that the dissociation Cumming discovered is really between early and late reports rather than between button pressing and verbal reports. Lachter et al (2000) and Lachter and Durgin (1999) found an advantage for early (speeded) reports over slower reports in a meta-contrast masking study, but only when the masking of the stimulus was strong. Information about the target stimulus appears to be briefly available for report before being 'over-written' by the appearance of the mask.

⁶ It might be thought that Marcel's subjects were merely failing to give a positive response rather than giving a negative response. But Marcel reports that he got the same results when his subjects were instructed to make a motor response for negative trials.

that are obviously precluded to her (such as knitting). Other anosognosic patients behave in ways that indicate that they are aware of their deficit, yet when asked about their condition they resolutely deny that there is anything wrong with them. Still other patients give inconsistent self-evaluations: they may deny hemiplegia but admit that the affected limbs are 'lazy' or 'naughty'. Hemiplegic patients with anosognosia may claim that although they themselves are able to engage in certain tasks—such as climbing a ladder—that they are obviously incapable of, other people would be unable to engage in these activities if they were affected by the same impairment.

In each of the syndromes just reviewed subjects appear to have conscious states whose content are only locally available to consuming systems. As such, these phenomena provide *prima facie* evidence against the availability thesis. Of course, none of these syndromes provides knockdown objection to the availability thesis. I have suggested that these cases involve conscious states whose contents are only partially available to the subject's consuming systems, but it might be argued instead that they involve the high-level availability of non-conscious content. Perhaps the sort behavior of children in the Dimensional Change Card Sort task is guided by non-conscious representations of the pre-switch rule, and perhaps the visual states that drove the button-pressing (guessing) responses in Cumming's experiments were not conscious. Alternatively, the availability theorist might insist that these states are conscious, but that consciousness (phenomenality) is correlated not with global availability to the subject's consuming systems but with the availability to each of a certain set of the subject's consuming systems. Both lines of response have something to be said for them, but I lack the space to evaluate them here. My goal here is to draw attention to the kinds of cases that put pressure on the availability thesis, and to argue that the thesis is far less straightforward than it is often thought to be.

5. The unity thesis

The third of my unity theses concerns the relation between subject unity and phenomenal unity. Consider what it is like to be a subject of experience. One typically enjoys a variety of perceptual, cognitive, emotional and agentive experiences, but no matter how numerous, varied, or complex these experiences

are, they occur as parts (components, aspects) of a single, global phenomenal state. It is the content of this global state of consciousness that determines what it is like to be you right now.

Arguably, this phenomenal unity extends beyond normal waking phenomenology to include even non-standard forms of consciousness, such as those experienced while dreaming and in states of delirium. Building on this thought, one might argue that for any subject of experience, there will be a global phenomenal state that subsumes each of the experiences that the subject in question has at that time. We can call this proposal the *unity thesis* (Bayne & Chalmers 2003). If the unity thesis is right, then there is a deep and intimate connection between subject unity and phenomenal unity.

Engaging in a full evaluation of the unity thesis is obviously beyond the scope of this chapter. Here, I will be content to argue against an influential objection to the unity thesis—the ‘split-brain objection’.

The basic split-brain syndrome is produced by sectioning the corpus callosum—the bundle of fibers that serve as the primary channel of communication between the two hemispheres of the brain. This procedure has little impact on cognitive function in everyday life, but careful research has revealed a complex array of deficits—and the occasional benefit—in the split-brain (Bogen 1993; Corballis 1995; Gazzaniga 2000, 2005).

The standard methodology for studying perception in the split-brain involves projecting distinct stimuli to the patient’s two hemispheres.⁷ Consider a typical split-brain patient (SB). The word ‘key-ring’ might be presented so that ‘key’ falls within the SB’s left visual field and ‘ring’ falls within SB’s right visual field. The contralateral structure of the visual system ensures that stimuli projected to the left visual field are processed in the right hemisphere and vice-versa. When asked to report what she sees SB will say only that she sees the word ‘ring’, yet, with her left hand SB may select a picture of a key, ignoring pictures of a ring and a key-ring.

⁷ See Gazzaniga (2000) for a useful introduction to the techniques used to study the split-brain.

It is widely assumed that this behavior demonstrates that SB has two simultaneous streams of consciousness, at least within experimental contexts (Davis 1997; Gazzaniga & LeDoux 1978; Marks 1981; Moor 1982; Puccetti 1981; Tye 2003). Two sorts of arguments are given for this conclusion. Some theorists appeal to the SB's apparent *representational* disunity: SB appears to have conscious representations of the words 'key' and 'ring', but no conscious representation of the word 'key-ring'. Other theorists appeal to SB's apparent *access* disunity: SB's representation of 'key' is available to some consuming systems, and her representation of 'ring' is available to other consuming systems, none (or only very few) of the subject's consuming systems has access to both representations.

I think it is reasonable to grant that if split-brain patients suffer from the kinds of representational and access disunities just outlined then there is a strong case against the unity thesis. But there is little reason to think that split-brain subjects do suffer from these representational and behavioral disunities. A careful examination of the evidence suggests that conscious perception in the split-brain subject may alternate between their hemispheres, rather than each hemisphere supporting its own stream of consciousness. The main evidence for this claim comes from research conducted by Levy and collaborators involving chimeric stimuli, that is, stimuli created by conjoining similar stimuli at the vertical midline (Levy 1977; 1990; Levy & Trevarthen 1976; Levy, Trevarthen & Sperry 1972; Trevarthen 1974). Since each hemisphere received a different stimulus one would expect the subject to produce conflicting motor response if representations of both stimuli were processed up to conscious levels. For example, one would expect the patient to verbally identify the stimulus as face A while using his or her left-hand to point to face B. In fact, such responses were vanishingly rare:

For all patients examined, and for tasks including the perception of faces, nonsense shapes, picture of common objects, patterns of Xs and squares, words, word meaning, phonetic images of rhyming pictures, and outline drawings to be matched to colors, patients gave one response on the vast majority of competitive trials. Further, the nonresponding hemisphere gave

no evidence that it had any perception at all. Thus, if the right hemisphere responded there was no indication, by words or facial expression, that the left hemisphere had any argument with the choice made, and, similarly, if the left hemisphere responded, no behavior on the part of the patient suggested a disagreement by the right hemisphere. (Levy 1990: 235; see also Levy 1977)

Levy and Trevarthen found that asking patients to match chimeric stimuli based on their visual appearance typically favoured the figure presented to their LVF (implicating the right hemisphere) whereas instructing the subjects to match the stimuli based on their function typically favoured the figure presented to their RVF (implicating the left hemisphere) (Levy & Trevarthen 1976). But at no point did both hemispheres appear to sustain simultaneous conscious perceptions. This reallocation of attention not only changes the content of the patient's experience, it also changes the consuming systems to which the contents of their experience is available.

How then might we explain the anecdotal 'key-ring' evidence? Arguably, the same sort of explanation applies to the 'key-ring' behavior that is anecdotally reported. Although it is sometimes implied that inconsistent behaviours such as this are carried out simultaneously (see e.g. Gazzaniga & LeDoux 1978), I am unaware of any quantitative support for such claims. Further, it is possible that the conscious states underlying simultaneous behaviours might not be simultaneous. Either way, it is possible that the state of consciousness underlying this behaviour involves an experience of 'ring' followed by an experience of 'key' (or vice-versa), rather than simultaneous (but disunified) experiences of 'key' and 'ring'.

The findings that Levy and Trevarthen report indicate that there is no straightforward 'split-brain' objection to the unity thesis. Split-brain subjects do, of course, suffer from certain kinds of disunities in consciousness, but there is little reason to think that they have two simultaneous streams of consciousness at once. Of course, there are other challenges to the unity thesis, but it is beyond the scope of this chapter to examine them.

6. Conclusion

In this chapter I have examined three conceptions of the unity of consciousness—the consistency thesis, the availability thesis, and the unity thesis. Each thesis attempts to capture a central respect in which consciousness is, or at least appears to be, unified. I have argued that the consistency and availability theses face serious objections, and I have suggested that an important objection to the unity thesis might not be as potent as it is often assumed to be. But rather than exploring any one conception of the unity of consciousness in detail, my primary goal has been to present a framework of the unity of consciousness which can be used both for understanding various syndromes and for evaluating accounts of consciousness. Filling in this framework is a task that must be left for another occasion.⁸

References

- Allport, A. (1988). "What concept of consciousness?" In A.J. Marcel and E. Bisiach (eds.), *Consciousness in Contemporary Science*. Oxford: Clarendon.
- Baars, B. (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Baars, B. (1993). "How does a serial, integrated and very limited stream of consciousness emerge from a nervous system that is mostly unconscious, distributed, and of enormous capacity?" In G. Bock and J. Marsh (eds.), *CIBA Symposium on Experimental and Theoretical Studies of Consciousness*. London: Wiley, 282-290.
- Baars, B. (2002). "The conscious access hypothesis: origins and recent evidence". *Trends in Cognitive Sciences*, **6/1**: 47-52.
- Bayne, T. and D. Chalmers (2003). "What is the unity of consciousness?" In A. Cleeremans (ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*. Oxford: Oxford University Press, 23-58.

⁸ Many thanks to Jillian Craigie and Neil Levy for their helpful comments on an earlier version of this chapter.

- Bisiach, E. and A. Berti (1995). "Consciousness in dyschiria". In M. Gazzaniga (ed.), *The Cognitive Neurosciences*. Cambridge, MA: MIT Press, 1331-1340.
- Bogen, J.E. (1993). "The callosal syndromes". In K.H. Heilman and E. Valenstein (eds.), *Clinical Neuropsychology*. Oxford: Oxford University Press, 337-407.
- Corballis, M. (1995). "Visual integration in the split-brain". *Neuropsychologia*, **33/8**: 937-959.
- Crane, T. (1988). "The waterfall illusion". *Analysis* **48**: 142-147.
- Dainton, B. (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.
- Davis, L. (1997). "Cerebral hemispheres". *Philosophical Studies* **87**: 207-22.
- Dehaene, S and L. Naccache (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework'. *Cognition* **79**: 1-37.
- Dennett, D. (2001). "Are we explaining consciousness yet?" *Cognition* **79**: 221-37.
- Fleminger, S. (2002). "Remembering delirium". *British Journal of Psychiatry* **180**: 4-5.
- Frisby, J.P. (1979). *Seeing*. Oxford: Oxford University Press.
- Gazzaniga, M.S. (2000). "Cerebral specialization and interhemispheric communication. Does the corpus callosum enable the human condition?" *Brain* **123**: 1293-1336.
- Gazzaniga, M.S. (2005). "Forty-five years of split-brain research and still going strong". *Nature Reviews Neuroscience* **6**: 653-59.
- Gazzaniga, M.S. and J. LeDoux (1978). *The Integrated Mind*. New York: Plenum Press.
- Gill, D. and R. Mayou (2000). "Delirium". In *The New Oxford Textbook of Psychiatry*, Oxford: Oxford University Press, 382-387.
- Hurley, S. (1998). *Consciousness in Action*. Cambridge, MA: Harvard University Press.

- Hurley, S. (2000). "Clarifications: Responses to Kobes and Kinsbourne". *Mind and Language* **15**: 556-561.
- Kohler, I. (1961). "Experiments with goggles". *Scientific American* **206**: 62-86.
- Lachter, J. and F. Durgin (1999). "Metacontrast masking functions: A question of speed"? *Journal of Experimental Psychology: Human Perception and Performance* **25**: 936-47.
- Lachter, J., F. Durgin, and T. Washington (2000). "Disappearing percepts: Evidence for retention failure in metacontrast masking". *Visual Cognition*, **7**: 269-79.
- Levy, J. (1977). "Manifestations and implications of shifting hemi-inattention in commissurotomy patients". *Advances in Neurology* **18**: 83-92.
- Levy, J. (1990). "Regulation and generation of perception in the asymmetric brain". In C. Trevarthen (ed.) *Brain Circuits and Functions of the Mind: Essays in Honour of Roger W. Sperry*. Cambridge: Cambridge University Press, 231-48.
- Levy, J. & C. Trevarthen (1976). "Metacontrol of hemispheric function in human split-brain patients". *Journal of Experimental Psychology: Human Perception and Performance* **2**: 299-312.
- Levy, J., C. Trevarthen and R.W. Sperry (1972). "Perception of bilateral chimeric figures following hemispheric deconnexion". *Brain* **95**: 61-78.
- Lipowski, Z.J. (1990). *Delirium: Acute Confusional States*. New York: Oxford University Press.
- Lockwood, M. (1989). *Mind, Brain and the Quantum: The Compound "I"*. Oxford: Blackwell.
- Logothetis, N.K., D.A. Leopold, and D.L. Sheinberg (2003). "Neural mechanisms of perceptual organization". In N. Osaka (ed.), *Neural basis of consciousness. Advances in Consciousness Research: Vol. 49*. Amsterdam: John Benjamins, 87-103.

- Marcel, A. (1993). "Slippage in the unity of consciousness". In G.R. Bock and J. Marsh (eds.), *Experimental and Theoretical Studies of Consciousness*. Chichester: John Wiley and Sons, 168-179.
- Marcel, A. (1994). "What is relevant to the unity of consciousness?" In C. Peacocke (ed.) *Objectivity, Simulation, and the Unity of Consciousness*. Oxford: Oxford University Press, 79-88.
- Marcel, A., R. Tegnér, and I. Nimmo-Smith (2004). "Anosognosia for plegia: Specificity, extension, partiality and disunity of bodily awareness". *Cortex* **40**: 19-40.
- Marks, C. (1981). *Commissurotomy, Consciousness and Unity of Mind*. Cambridge, MA: MIT Press.
- Moor, J. (1982). "Split-brains and atomic persons". *Philosophy of Science* **49**: 91-106.
- Puccetti, R. (1981). "The case for mental duality: Evidence from split-brain data and other considerations". *Behavioral and Brain Sciences* **4**: 93-123.
- Rubin, N. (2003). "Binocular rivalry and perceptual multi-stability". *Trends in Neurosciences* **26**: 289-291.
- Shoemaker, S. (2003). "Consciousness and co-consciousness". In A. Cleeremans (ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*. Oxford: Oxford University Press, 59-71.
- Taylor, J.G. (1962). *The Behavioural Basis of Perception*. New Haven, CN: Yale University Press.
- Tononi, G. and G.M. Edelman (1998). "Consciousness and complexity". *Science*, **282**: 1846-1851.
- Trevarthen, C. (1974). "Functional relations of disconnected hemispheres with the brain stem, and with each other: monkey and man". In M. Kinsbourne and W. Lynn Smith (eds.) *Hemispheric Disconnection and Cerebral Function*. Springfield, Il: Charles C. Thomas, 187-207.
- Tye, M. (2003). *Consciousness and Persons*. Cambridge, MA: MIT Press.

Vuilleumier, P. (2004). "Anosognosia: The neurology of beliefs and uncertainties". *Cortex* **40**: 9-17.

Zelazo, P.D. (1996). "An age-related dissociation between knowing rules and using them". *Cognitive Development* **11**: 37-63.

Zelazo, P.D. (2004). "The development of conscious control in childhood". *Trends in Cognitive Sciences* **8/1**: 12-17.